

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to the Department of Defense, Executive Service Directorate (0704-0188). Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.</p>					
1. REPORT DATE (DD-MM-YYYY) 12-28-2007		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) June 1, 2002 to Sept 30, 2007	
4. TITLE AND SUBTITLE COOPERATIVE NETWORKED CONTROL OF DYNAMICAL PEER-TO-PEER VEHICLE SYSTEMS				5a. CONTRACT NUMBER F49620-02-1-0325	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) G.E. Dullerud (PI), F. Bullo, E. Feron, E. Frazzoli, P.R. Kumar, S. Lall, D. Liberzon, N. Lynch, J.C. Mitchell, S. Mitter, E. Modiano, P. Parrilo, B. Reznick, M. Viswanathan				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Univesity of Illinois, Office of Sponsored Programs and Research, 1901 S. First Street, Champaign, Illinois 61820-1046 (with subcontractors MIT, Stanford University, University of California).				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Office of Scientific Research, 875 N Randolph St, Arlington, VA 22203-1768 <i>Scott Weeks/NL</i>				10. SPONSOR/MONITOR'S ACRONYM(S) AFOSR	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release. Distribution is unlimited AFRL-SR-AR-TR-08-0025					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT The goal of this MURI center was the development of a rigorous theoretical foundation, and scalable analytical tools and paradigms for construction of networked control for large numbers of autonomous and semi-autonomous air vehicles. The research is specifically aimed at the critical reliability and performance issues facing autonomous vehicle systems which operate in highly uncertain environments, and enables the vehicles to form teams, manage information, and coordinate operations including deployment, task allocation and search. The program produced both the fundamental theory necessary to allow systematic performance analysis, verification and validation of such systems, as well as algorithms for implementation, and design software. Specifically, advances have been made in dynamic deployment and task allocation; verification and hybrid systems; and information management for cooperative control. The activity of the program has had a significant impact on the understanding and designing of large-scale cooperative UAV systems, providing the long-term basis for major new capability, and will more generally enable systematic construction of large-scale robust real-time distributed systems. The members of the team have won numerous major research awards and recognitions during the course of the project.					
15. SUBJECT TERMS cooperative control, multi-vehicle deployment algorithms, dynamic routing of autonomous vehicles, validation and verification of complex systems, decentralized control, delay aware wireless networks					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (Include area code)

**FINAL REPORT
COOPERATIVE NETWORKED CONTROL OF
DYNAMICAL PEER-TO-PEER VEHICLE SYSTEMS**

F49620-02-1-0325

Geir E. Dullerud
Mechanical Engineering, University of Illinois

Francesco Bullo
Mechanical and Environmental
Engineering
UCSB

Eric Feron
Aerospace Engineering
Georgia Institute of Technology

Emilio Frazzoli
Aeronautics and Astronautics
MIT

P.R. Kumar
Electrical and Computer
Engineering, University of Illinois

Sanjay Lall
Aeronautics and Astronautics
Stanford University

Daniel Liberzon
Electrical and Computer
Engineering, University of Illinois

Nancy A. Lynch
Electrical Engineering and
Computer Science, MIT

John C. Mitchell
Computer Science
Stanford University

Sanjoy K. Mitter
Electrical Engineering and
Computer Science, MIT

Eytan Modiano
Aeronautics and Astronautics
MIT

Pablo Parrilo
Electrical Engineering and
Computer Science, MIT

Bruce Reznick
Mathematics
University of Illinois

Mahesh Viswanathan
Computer Science
University of Illinois

20080118118

Contents

1	EXECUTIVE SUMMARY	4
2	MULTI-VEHICLE PEER-TO-PEER DEPLOYMENT SCENARIOS	6
2.1	Deployment and Coverage Control for Multi-Vehicle Networks	6
2.1.1	Area-coverage deployment	7
2.1.2	Deployment for maximum detection likelihood	8
2.2	Visibility-based Deployment in Nonconvex Environments	9
3	DYNAMIC MULTI-VEHICLE ROUTING	13
3.1	Dynamic Traveling Repairperson Problem: Decentralized Policies	14
3.2	The Dubins Traveling Salesperson Problem (DTSP)	14
3.2.1	From the Euclidean to the Dubins TSP	15
3.2.2	The DTSP in the worst case	15
3.2.3	The Stochastic DTSP	16
3.3	The DTRP for Dubins vehicles under heavy load	17
3.4	The TSP and DTRP for other vehicle models	18
3.5	Scheduling Multiple Vehicles Dynamically and Bandit Problems	18
3.6	Risk-Sensitive Multi-Agent Systems	19
4	CONTROL WITH NETWORK GRAPH CONSTRAINTS	19
4.1	Multiplayer Games	20
4.1.1	Polynomial and semialgebraic games	20
4.2	Decentralized Control	22
4.3	Distributed Control of Heterogenous Systems	23
4.4	Decentralized Control of Markov Processes	24
5	VERIFICATION AND VALIDATION FOR COMPLEX SYSTEMS	26
5.1	Input-Output Automata Modeling Frameworks and Methodology	26
5.1.1	Mathematical foundations	26
5.1.2	Computer-assisted tools and formal techniques	28
5.1.3	Vehicle-based examples and applications	29
5.2	Learning and Randomization	31
5.2.1	Learning to verify	31
5.2.2	Randomization in verification	32
5.3	Decidability Results on Discrete and Hybrid Systems	33
5.4	Switched Systems	33
5.5	Markovian Jump Systems: Uniform Performance	35
5.6	Certificates for Nonlinear Dynamics	35
5.6.1	Nonlinear system theory	35
5.6.2	Polynomial computation of invariant sets	36
5.6.3	Semialgebraic fundamentals	37
6	COMMUNICATIONS FOR NETWORKED CONTROL	38
6.1	Control-oriented Information Theory	38
6.2	Control with Limited Information: Deterministic Systems	39
6.3	Delay Aware Wireless Networks	41
6.3.1	Fundamental delay bounds	41
6.3.2	Transmission scheduling schemes for time-critical data	41

6.3.3	Wardrop routing	41
6.4	Synchronization of Clocks in Wireless Systems	42
6.5	Performance in Mobile Wireless Systems	43
6.5.1	Optimal flow control scheme for maximizing network throughput and utility .	43
6.5.2	Joint routing and power allocation for wireless networks	43
6.5.3	Randomized algorithms for distributed network control	43
6.5.4	Maximizing network throughput via partitioning	44
6.5.5	Transmission scheduling for MIMO channels	44
6.6	Increasing Reliability in Cooperative Routing	44
7	PUBLICATIONS	45
8	APPENDICES	57
A	Awards and Honors	57
B	Transitions	58
C	Personnel Supported (list cumulative)	58

1 EXECUTIVE SUMMARY

This MURI center was focused on the development of a rigorous theoretical foundation, and scalable analytical tools and paradigms, for construction of networked control for large numbers of autonomous and semi-autonomous air vehicles. The results have been a substantial body of research accomplishments that are currently having an impact, and can be expected to provide the long-term foundation and organizing principles for the development of cooperative multi-vehicle capabilities. This research has specifically targeted the critical reliability and performance issues facing autonomous vehicle systems operating in highly uncertain environments, and enables vehicles to form teams, manage information, and coordinate operations including deployment, task allocation and search. The program concentrated on both the fundamental theory necessary to allow systematic performance analysis, verification and validation of such systems, as well as the development and implementation of algorithms and software. Cooperative control of multi-vehicle systems requires fundamental and coordinated advances in algorithms for control, communications, and computing to develop systems that are verifiably robust. The research of this MURI center provides required core algorithms and internal software methodologies, and several transitions of the research have now occurred (see Appendix B). The activity of the program has impacted the understanding of large-scale cooperative unmanned aerial vehicle (UAV) systems, providing a basis for major new technical and operational capabilities, and will more generally enable systematic construction of large-scale, robust, real-time distributed systems. This MURI program has influenced the directions pursued by the cooperative control research community at large, and the members of the team have won numerous major research awards and recognitions during the course of the project (see Appendix A); a cumulative list of personnel supported is provided in Appendix C.

This project has significant research accomplishments in three primary areas: cooperative deployment and task allocation algorithms; rigorous verification and validation for complex systems; and information management protocols for cooperative vehicle control. The specific accomplishments have direct application to large-scale multi-UAV systems and we now provide an executive overview of the research achievements.

Deployment and Task Allocation. A general class of algorithms of primary importance that has been developed in this program is that focused on independent deployment of air vehicles for surveillance. The associated algorithms run in real-time on board each vehicle to route them to optimal locations, and are strongly decentralized, coordinating between vehicles, enabling them to efficiently deploy throughout a geographic region. More generally, also considered have been the related goals of coordinated scouting, rendezvous and task or target allocation. A significant feature of this research, as with all the research of the project, is provable guarantees for the performance of the algorithms developed.

Significant research accomplishments have been made on deployment and rendezvous in convex environments. More recent achievements have focused on non-convex situations, such as encountered in urban settings or more expansively any environment with obstacles. A specific deployment problem considered is that where multiple vehicles have the objective of providing full visibility of the environment, and must do so using a line-of-sight wireless communication scheme; this is a so-called visibility-based deployment problem. The problem is closely related to many surveillance and pursuer-evader problems, and the results obtained are adaptive, distributed, asynchronous, and verifiably correct, and provide very sharp sufficient estimates for the number of vehicles needed to guarantee that the task be achieved.

Deployment problems are mainly motivated by either surveillance applications, or by scenarios where vehicles are optimally positioned to service critical, but infrequently occurring, targets. As a distinct research area, algorithms have been successfully developed for scenarios involving multiple vehicles and targets, where targets appear frequently and dynamically, modelled by spatio-temporal Poisson processes. The case where the vehicles have dynamical constraints is considered, and the

objective is to service the stochastically appearing targets in such a way as to minimize the wait times to servicing. The research is closely related to the so-called Travelling Repairperson Problem (TRP), and has produced the best available algorithm for this problem that guarantees performance. Also, in very novel work discrete resource allocations problems have been considered, using a decentralized Markov decision process model, where possible targets appear at a fixed known set of points. For these classes of computationally hard problems we have developed semidefinite relaxations that yield polynomial-time algorithms which can find control policies that have performance within a small constant factor of the optimum.

Significant advances have been made to optimal target servicing involving vehicles that have nonholonomic constraints. The motivation for the issues considered is that air vehicles, and in particular UAVs, have limited turning capability and cannot directly reverse their motion. These features have been modelled using a Dubins' vehicle model, and very strong results have been achieved for minimum-time motion planning and routing problems for such vehicles which are constrained to move along planar paths of bounded curvature, without reversing direction. Various additional scenarios of *dynamic vehicle routing* problems for a group of autonomous vehicles have been studied, and using novel concepts and algorithms have been solved.

Verification and hybrid systems. The program has produced significant advances in the theory of hybrid input-output automata (HIOA) and the resulting verification techniques; these techniques enable off-line automatic verification and validation of safety and liveness of cooperative control algorithms, such as those discussed in the preceding paragraphs. This HIOA approach combines ideas from control theory and techniques from input-output automata theory to achieve verification of cooperative control algorithms. The methods allow complex proofs to be broken down into manageable pieces, and provide rigorous techniques for establishing both safety and liveness properties. The research has also extended the IOA ideas to probabilistic IOA, which are specifically needed to allow for systematic and automated reasoning about algorithms that have stochastic components. A monograph has been published during the project that summarizes many of the above results. The software company VeroModo was founded during the course of the project to make the above techniques commercially available in computer-assisted form.

The concept of Virtual Node Layers (VNLayers) for implementing algorithms in distributed vehicle networks has been developed. VNLayers are abstraction layers for simplifying the task of programming mobile networks, in much the same way virtual machines and high-level programming languages simplify the task of programming ordinary stand-alone computers. A VNLayer masks dynamic, unpredictable behavior on part of the underlying sensor network, including node and communication failures, joining, leaving, and mobility. It produces more robust, easier-to-program, higher-level network abstraction. Many different varieties of Virtual Node can be defined, satisfying different assumptions regarding the types of operations they may perform, their knowledge of geography and time, their control over the timing of their steps, whether they are stationary or mobile (and if mobile, how they are allowed to move), and their failure modes. Several application scenarios have been explicitly studied illustrating the power of this approach.

Research on formal verification of systems using model checking, a direct computer-assisted approach, with very large state spaces and uncertainty has been carried out with several significant accomplishments; the uncertainty in these systems being modelled by probabilistic transitions. Verifying these systems automatically is a difficult problem, and doing so has been tackled using two paradigmatic approaches: statistical sampling, and machine learning. A learning based model checking algorithm to verify safety properties of general infinite state systems has also been developed. These ideas have also been extended to model checking general branching time properties expressed in CTL. Also developed have been learning algorithms for boolean programs, which are sequential, recursive programs over boolean variables. The software tool Vesta has been developed to make this research directly available to practitioners.

Information management for cooperative control. This program has created an information theory for closed-loop systems, which allows for analysis of the performance implications of packet networks on control systems. This research is very important because traditional information theory does not take account of delay, a consideration which is crucial in multi-vehicle and more generally control systems. This new work is fundamental, showing that the Shannon notion of channel capacity is ill-suited to feedback tasks, and is having significant influence and impact on the multi-vehicle research community.

This program has made significant contributions to the design of communication protocols with the robustness and performance constraints required for cooperative multi-vehicle systems. This includes cooperative routing schemes which take advantage of network layer diversity, and delay adaptation, to increase network reliability over a wireless network with fading channels. The developed protocols successfully overcome many of the architectural challenges involved with the software implementation of multi-path, delay feedback based, probabilistic routing algorithms. Several additional accomplishments have also been made in the area of performance of mobile wireless networks.

A more detailed account of the technical accomplishments of the Center now follows.

2 MULTI-VEHICLE PEER-TO-PEER DEPLOYMENT SCENARIOS

One fundamental capability of future networks of autonomous and semi-autonomous vehicles will be the ability to perform *spatially-distributed sensing tasks* including coverage, surveillance, exploration, target detection, and search. These future mobile and tunable sensor networks will be able to adapt to changing environments and dynamic situations, will provide guaranteed fault-tolerant quality of service, and will operate via limited-bandwidth ad-hoc communication links. To achieve these desirable capabilities, the major objective of this project is the design of multi-vehicle coordination algorithms that are distributed, asynchronous, adaptive, and verifiably correct.

In this section we describe two types of deployment problems: the first problem is setup in a convex environment and contains a notion of optimal positioning in locational optimization — the second problem is housed in a nonconvex environment and generalized to the distributed setting the classic art-gallery theorem in computational geometry. The resulting coordination protocols are based either on distributed descent algorithms and on aggregate utility functions that encode optimal coverage and sensing policies, or on a combination of distributed information gathering and useful geometric structures. From a broader perspective, the proposed approach unifies concepts and methods from systems theory, distributed algorithms, and algorithmic robotics.

2.1 Deployment and Coverage Control for Multi-Vehicle Networks

Here the results obtained in [34, 35, 88, 89] are outlined. The objective of this research is to develop a complete set of primitives for deployment and motion coordination in multi-vehicle networks. Multi-vehicle coordination is dealt with in a comprehensive fashion, developing fundamental modeling tools, metrics for performance analysis, and algorithmic design. In particular, it is of central importance to design algorithms that will gently scale with the number of vehicles and devices present in the network. The problems of optimal deployment and coverage are tackled in numerous variations. This class of problem is very broad and the features of specific formulations vary drastically with the underlying physical assumptions. Critical parameters include:

1. the environment of interest can be two or three dimensional, known or unknown, uniform or nonuniform (e.g., portions of the environment might be of greater interest), stationary or non-stationary (e.g., boundaries and nonuniformity may depend on time);

2. the deployment objectives can vary depending on the ultimate network objective: **examples** include search and exploration, target detection, localization and tracking, wireless communication coverage, environmental monitoring;
3. the communication and sensing characteristics of individual vehicle can be uniform or **heterogeneous** (e.g., antennas and sensors can be directional or omni-directional), the vehicle **mobility** and dynamics can vary drastically.

Technical approach to deployment and motion coordination

In what follows we illustrate the results we have obtained in some aspects of this broad **theme**. The following performance metrics and coordination algorithms are meant to illustrate the **proposed** approach and not to restrict our research objectives to any specific setting. The general “**bottom-up**” approach is to design basic behaviors, formalize the resulting network model through **nonlinear** and hybrid systems theory, and prove converge correctness via Lyapunov and invariant theory.

We discuss mainly deployment problems. Here, let Q be a region in \mathbb{R}^3 and let $\|\cdot\|$ be the Euclidean distance. Let $P = (p_1, \dots, p_n)$ be the location of n agents, each moving in the environment Q .

2.1.1 Area-coverage deployment

Let $\phi : Q \rightarrow \mathbb{R}_+$ play the role of a distribution density function; i.e., ϕ measures how many **users** of the communication channel are present, or how important it is to cover a certain region in the environment Q . Because of noise and loss of resolution, the sensing or communication performance at point $q \in Q$ taken from i th agent at the position p_i degrades with the distance $\|q - p_i\|$; we **describe** this degradation with a monotone (decreasing) function $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$. In other words, $f(\|q - p_i\|)$ is a point-wise quantitative assessment of how poor the sensing/communication performance is. Since typical agents have limited-range footprint, it is realistic to assume that $f(\|q - p_i\|)$ is **constant** (equally poor) outside the sphere $B_r(p_i)$ centered at p_i of radius r . As specific example, we let $f(\|q - p_i\|)$ equals 1 if q is inside the sphere $B_r(p_i)$ and 0 otherwise. This performance function leads to the following interpretation: the agent i provides equally good sensing/communication coverage over all points in its sphere of influence.

In a first approximation, let us assume that each individual agent is uniquely responsible for wireless coverage and measurements taken over a region to be determined. Let $\mathcal{W} = \{W_1, \dots, W_n\}$ be a collection of n regions with disjoint interiors whose union is Q ; we call \mathcal{W} a partition of Q and W_i the dominance region of agent i . Consider the coverage performance metric $\mathcal{H}(P, \mathcal{W}) = \sum_{i=1}^n \int_{W_i} f(\|q - p_i\|) d\phi(q)$. The function \mathcal{H} is to be maximized with respect to the agents location P and to the assignment of the dominance regions \mathcal{W} . One can easily see that, at fixed locations (p_1, \dots, p_n) , the optimal partition is the Voronoi partition $\mathcal{V}(P) = \{V_1, \dots, V_n\}$ defined by $V_i = \{q \in Q \mid \|q - p_i\| \leq \|q - p_j\|, \forall j \neq i\}$. Therefore, an equivalent expression of optimal coverage is $\mathcal{H}(P, \mathcal{V}(P)) = E[\max_{i \in \{1, \dots, n\}} f(\|q - p_i\|)]$. Remarkably, one can show that

$$\frac{\partial \mathcal{H}}{\partial p_i}(P, \mathcal{V}(P)) = \int_{V_i \cap B_r(p_i)} \frac{\partial}{\partial p_i} f(\|q - p_i\|) d\phi(q),$$

and deduce the following critical property: the gradient of \mathcal{H} is decentralized in the sense that it can be computed with information localized to each individual sphere of influence and Voronoi cell. Closed-form expressions for this partial derivative can be computed under various assumptions on the shape of f .

[Algorithm design] Finally, we design a deployment algorithm under the assumption that each agent location obeys a first order dynamical behavior described by $\dot{p}_i = u_i$. Set $u_i = \frac{\partial \mathcal{H}}{\partial p_i}(P, \mathcal{V}(P)) - p_i$, where $\mathcal{V}(P) = \{V_1, \dots, V_n\}$ is continuously updated in a decentralized computation. This closed-loop system is a gradient flow for the cost function \mathcal{H} so that performance is indeed locally, continuously optimized. The coverage optimization function \mathcal{H} is a Lyapunov function and the group of

mobile agents is guaranteed to converge to a local maximum of \mathcal{H} . Fig. 1 illustrates the performance of this coordination algorithm when Q is a 2D convex polygon.




Figure 1: Area-coverage deployment for 16 agents; the region of interest is characterized by a density function equal to the sum of Gaussians. The left (resp. right) figure contains the contour plot of the density function, the initial (resp. final) position of the agents, the agents' sphere of influence and Voronoi partitions. The central figure illustrates the joint motion.

2.1.2 Deployment for maximum detection likelihood

Next, we describe a second formulation of deployment with a different network objective. We consider n mobile devices equipped with acoustic sensors attempting to detect, identify, and localize a sound-source (we could similarly envision antennas detecting RF signals, or chemical sensors localizing a source). For a variety of criteria, when the source emits a known signal and the noise is Gaussian, we know that (1) the optimal detection algorithm involves a matched filter; (2) detection performance is a function of signal-to-noise-ratio; and, in turn, (3) signal-to-noise ratio is inversely proportional to the sensor-source distance. The goal is to deploy the agents and optimize their location to maximize the detection probability.

Recall that, the circumcircle of a given polygon is the smallest circle enclosing the polygon; circumradius and circumcenter and radius and center of the circumcircle, respectively. Given this notion, we introduce the following simple algorithm. If each agent moves toward the circumcenter of its Voronoi cell, then, as a function of time, the detection likelihood is inversely proportional to the circumradius of each agent's Voronoi cell, and the detection likelihood is monotonically increasing, see Fig. 2.




Figure 2: Deployment of 16 agents for maximum likelihood detection. The left (resp. right) figure contains the initial (resp. final) position of the agents, and the Voronoi partitions and circumcircles of each agent.

The fundamental reason this behavior is correct is the existence of an appropriate Lyapunov function, with respect to which the given behavior is dissipative. It turns out that, as a function of

the agents' position, an appropriate cost function is the maximum of the radiuses of disks centered at each agent's position and covering each agent Voronoi cell.

2.2 Visibility-based Deployment in Nonconvex Environments

In this section we review the results obtained in [48, 49, 50] on a deployment problem for robots in nonconvex environments. We consider a group of robotic agents modeled as point masses, moving in a simple nonconvex polygonal environment, Q . Each agent has a unique identifier **UID**, say i . Let p_i refer to the position of agent i . Each agent is equipped with an omnidirectional line-of-sight range sensor. Thus, the agent can sense its star-shaped visibility set $V(p_i)$. It can communicate with any other agent within line-of-sight.

Asynchronous networks of visually-guided agents

Each agent has access to some memory \mathcal{M}_i . An agent i can broadcast its UID together with its memory contents to all agents inside its communication region. Such a broadcast is denoted by $\text{BROADCAST}_i(i, \mathcal{M}_i)$. It can also receive broadcasts from other agents. We also assume that there is a bounded time delay, $\delta > 0$, between a broadcast and the corresponding reception. Each agent repeatedly performs the following sequence of actions between any two wake-up instants, say instants T_l^i and T_{l+1}^i for agent i :

1. **SPEAK**, that is, send a **BROADCAST** repeatedly at times $T_l^i + k\delta$, where $k \in \mathbb{N}_0$, until it starts moving;
2. **LISTEN** during the time interval $[T_l^i, T_l^i + \lambda_l^i]$, for $\lambda_l^i \geq \delta$;
3. **PROCESS** and **LISTEN** during the interval $[T_l^i + \lambda_l^i, T_l^i + \lambda_l^i + \rho_l^i]$, for $\rho_l^i \geq 0$;
4. **MOVE** during the time interval $[T_l^i + \lambda_l^i + \rho_l^i, T_{l+1}^i]$.

Agent i , in the **MOVE** state, is capable of moving at any time $t \in [T_l^i + \lambda_l^i + \rho_l^i, T_{l+1}^i]$ according to

$$p_i(t + \Delta t) = p_i(t) + u_i,$$

where the control is bounded in magnitude by 1. The control action depends on time, on the memory $\mathcal{M}_i(t)$, and on the information obtained from communication and sensing. The subsequent wake-up instant T_{l+1}^i is the time when the agent stops performing the **MOVE** action and it is not predetermined.

Given this model, the goal is to design a provably correct discrete-time algorithm which ensures that the agents converge to locations such that each point of the environment is visible to at least one agent. This is the *visibility-based deployment problem* for visually-guided agents.

The vertex-induced partition and tree The first step in the algorithm is to partition the environment into star-shaped polygons and construct a graph to represent the partition.

We will need the following notation. If p is a point in the polygon Q , we let $V(p)$ denote the set of visible points from p . A set S is *star shaped* if there exists $p \in S$ such that $S \subset V(p)$; if S is star shaped, we let $\text{ker}(S)$ be its *kernel*, i.e., the set of points $k \in S$ such that $S \subset V(k)$. Finally, a *diagonal* of a nonconvex polygon Q is a segment inside Q connecting two vertices of Q (and therefore splitting Q into two polygons). A vertex of a polygon Q is *nonconvex* when the internal angle is greater than π .

Given a nonconvex polygon Q without holes and a vertex s of it, we compute a list $\{P_1, \dots, P_m\}$ of star-shaped polygons composing a partition of Q and a list $\{k_1, \dots, k_m\}$ of kernel points for each star-shaped polygon $\{P_1, \dots, P_m\}$. The computation of these quantities is discussed in the following algorithm and is illustrated in Figure 3.




Figure 3: Computation of the vertex-induced partition and tree in 5 steps

Vertex-Induced Partition and Tree Algorithm

- 1: set $k_1 = s$, and collect all vertices of Q visible from k_1 (see Fig. 3(1))
- 2: let P_1 be the polygon determined by these vertices (by definition $k_1 \in \ker(P_1)$) (see Fig. 3(2))
- 3: identify the edges of P_1 that are diagonals of Q ; call them *gaps*. For all gaps, place a new point, say k_2 , across the gap at a new vertex of Q such that k_2 sees the gap. (see Fig. 3(3))
- 4: repeat last three steps for new point k_2 , until all gaps have been crossed. (see Fig. 3(4))
- 5: define edges starting from s going to all kernel points and crossing all edges. (see Fig. 3(5))

We refer to the list $\{P_1, \dots, P_m\}$ computed in the algorithm as the *vertex-induced partition*. The algorithm computes not only the partition and a list of kernel points, but also a collection of edges connecting the kernel points. In other words, we also computed a directed graph, the *vertex-induced tree*, denoted by $\mathcal{G}_Q(s)$: the nodes of this directed graph are $\{k_1, \dots, k_m\}$ and an edge exists between any two vertices k_i, k_j if and only if $P_i \cap P_j$ is a diagonal of Q . Note that $k_1 = s$; we refer to this node as the root of $\mathcal{G}_Q(s)$. We now state some important properties of the vertex-induced tree.

Proposition 1. *Given a polygon Q without holes and a vertex s , the following statements hold:*

1. the directed graph $\mathcal{G}_Q(s)$ is a rooted tree;
2. the maximum number of nodes in the vertex-induced tree is less than or equal to $\lfloor \frac{n}{2} \rfloor$, where n is the number of vertices in Q .

It is clear from the construction of the vertex-induced tree that if we design a distributed algorithm to place agents on each node of the tree, then we will have solved the art-gallery deployment problem. In other words, the vertex-induced tree has been defined here in a *centralized* manner, but visually-guided agents will be able to explore it and compute it in an *incremental distributed* manner. This is the subject of the next subsections.

Local node-to-node navigation algorithms Note that by virtue of the constructions in the previous section, we have converted the original problem into a graph “navigation and deployment” problem. We now informally describe a distributed algorithm to cover the nodes of a given tree by means of local sensing and with limited communication and memory.

We begin with algorithms to plan paths between neighboring nodes of the vertex-induced tree. In a rooted tree, every neighbor of a node is either a child or the parent. Therefore, we present two simple informal descriptions.

Move-to-Child Algorithm

- 1: compute the mid-point of the gap between the node and the child
- 2: go to the mid-point
- 3: compute the nearest vertex from which the entire gap is visible and which is across the gap
- 4: go to that vertex

Move-to-Parent Algorithm

- 1: compute the shortest path between the node and the parent
- 2: go to the nonconvex vertex which is a part of the shortest path
- 3: from the nonconvex vertex, go to the vertex representing the parent node

Figure 4 shows paths between parents and children as computed by the previous two algorithms. It is easy to see that navigation is very simple if sufficient information is available to the agents. We address this aspect in the next subsection.




Figure 4: Left figure: a vertex-induced tree and partition in a prototypical floor-plan. Center and right figure: the planned paths “from node to parent” and “from node to children,” respectively.

Distributed information processing From the previous discussion we know that the following information must be available to an agent to properly navigate from node to node. If the node is executing the **Move-to-Child Algorithm**, then it needs to know what gap to visit, i.e., what child to visit. If the node is executing the **Move-to-Parent Algorithm**, then it needs to know where the parent node is located and what gap leads to it.

This geographic information is gathered and managed by the agents via the following **state** transition laws and communication protocols. At this time we make full use of the computation, communication and sensing abilities of visually-guided agents mentioned in the modeling discussion.

1. The memory content \mathcal{M} of each agent is a quadruple of points in Q labeled $\{p_{\text{parent}}, p_{\text{last}}, g_1, g_2\}$. All four values are initialized to the initial location of the agent. These values are broadcast together with the agent’s UID during the **SPEAK** action.

During run time, \mathcal{M} is updated to acquire and maintain the following meaning: p_{parent} is the parent kernel point to the current agent’s position, p_{last} is the last node visited by the agent, and (g_1, g_2) is the diagonal shared between the current cell and the parent cell, i.e., the gap toward the parent node. This is accomplished as follows:

2. After an agent moves from a kernel point k_{parent} to a child kernel point k_{child} through a gap described by two vertices v', v'' , its memory \mathcal{M} is updated as follows: $p_{\text{parent}} := k_{\text{parent}}$, $p_{\text{last}} = k_{\text{parent}}$ and $(g_1, g_2) := (v', v'')$.
3. After an agent moves from a kernel point k_{child} to the parent kernel point k_{parent} , its memory \mathcal{M} is updated as follows: first, $p_{\text{last}} := k_{\text{child}}$ and second, the agent acquires updated values of $\{p_{\text{parent}}, g_1, g_2\}$ by listening to incoming messages.

Global exploration and deployment algorithms At this time, we have all the elements necessary to present a global navigation algorithm that leads the agents to deploy themselves over the nodes of the vertex-induced tree.

Depth-First Navigation Algorithm

All agents are initially located at root s

During each PROCESS action, each agent executes:

- 1: Find maximum UID received during the LISTEN action
- 2: **If** this UID is less than its own UID
- 3: **then** stay at current kernel point
- 4: **else**
- 5: **If** there are no children of the present kernel point
- 6: **then** Move-to-Parent Algorithm towards p_{parent} via $\{g_1, g_2\}$
- 7: **else**
- 8: Order the children in a suitable way
- 9: **If** p_{last} in memory is the parent of the present node,
- then** Move-to-Child Algorithm towards the first child in the ordering
- 10: **If** the last node visited is a child that is not the last in the ordering,
- then** Move-to-Child Algorithm towards next child in the ordering
- 11: **If** the last node visited is a child that is the last in the ordering,
- then** Move-to-Parent Algorithm towards p_{parent} via $\{g_1, g_2\}$

Note that the instruction 5: through 11: in **Depth-First Navigation Algorithm** essentially amount to a depth-first graph search. Alternatively, it is fairly easy to design (1) breadth-first search algorithms, or (2) randomized graph search algorithms, where the nodes select their motion among equally likely children/parent decisions.

The following Figures 5 and 6 show the results of the simulations of the depth-first search and randomized search algorithms respectively. The nodes of the vertex-induced tree of the environment in the simulations are precisely the locations where the agents in Figure 5 are located at the end of the simulation. In Figure 6, there are more agents than the number of nodes in the vertex-induced tree. Hence, the extra agents keep exploring the graph without coming to rest.




Figure 5: From left to right, evolution of a network implementing depth-first search. The number of vertices of the environment is $n = 46$ and the number of agents is $N = 13 < \lfloor \frac{46}{3} \rfloor$. Each point of the environment is visible at the end of the simulation.

Run time analysis According to the **Move-to-Child Algorithm** and **Move-to-Parent Algorithm**, the path from a node to its parent is shorter than the path from the parent to the node. Given a polygon Q without holes and a vertex s , we define the following lengths: For each edge (k_j, k_i) of $\mathcal{G}_Q(s)$, let $d_{\text{ford}}(k_j, k_i)$ and $d_{\text{back}}(e_i)$ denote the path length from k_i to the parent k_j and from k_j to its child k_i , respectively. The forward and backward lengths of $\mathcal{G}_Q(s)$ are defined by

$$\mathcal{L}_{\text{ford}}(\mathcal{G}_Q(s)) = \sum_{e \in \text{edges of } \mathcal{G}_Q(s)} d_{\text{ford}}(e), \quad \mathcal{L}_{\text{back}}(\mathcal{G}_Q(s)) = \sum_{e \in \text{edges of } \mathcal{G}_Q(s)} d_{\text{back}}(e),$$




Figure 6: From left to right, evolution of a network implementing randomized search. While the polygon is the same as above and therefore the vertex-induced tree still has only 13 nodes, the number of agents is 15; after each node of the tree is populated, the 2 extra agents continue to explore the vertex-induced tree.

respectively. This discussion is now summarized in the following proposition.

Proposition 2 (Run Time Analysis). *Given a polygon without holes Q , assume that N visually-guided agents begin their motion from a vertex s of Q . Assume Q has n vertices and the vertex-induced tree $\mathcal{G}_Q(s)$ has $m \leq \lfloor n/2 \rfloor$ nodes. The following statements hold:*

1. *In finite time there will be at least one agent on $\min\{m, N\}$ nodes of $\mathcal{G}_Q(s)$.*
2. *If $N \geq \lfloor n/2 \rfloor$, then the art-gallery deployment problem is solved in finite time by the Depth-First Navigation Algorithm.*
3. *If there exist bounds λ_{\max} and ρ_{\max} such that $\lambda_i^l \ell_e \lambda_{\max}$ and $\rho_i^l \ell_e \rho_{\max}$ for all i and l , then*

$$t^* \ell_e \mathcal{T}_{\text{motion}} + \mathcal{T}_{\text{nodes}},$$

where $\mathcal{T}_{\text{motion}} \ell_e 2 \left(\mathcal{L}_{\text{ford}}(\mathcal{G}_Q(s)) + \mathcal{L}_{\text{back}}(\mathcal{G}_Q(s)) - \min_{e \in \text{edges of } \mathcal{G}_Q(s)} d_{\text{back}}(e) \right)$ and $\mathcal{T}_{\text{nodes}} \ell_e 2(m-1)(\lambda_{\max} + \rho_{\max})$.

3 DYNAMIC MULTI-VEHICLE ROUTING

Motion coordination strategies for groups of autonomous robots is an area of research with broad civilian and military applications. In this project we were concerned with the generation of efficient cooperative strategies for several mobile agents to move through a certain number of given target points, possibly avoiding obstacles or threats. Trajectory efficiency in these cases is understood in terms of cost for the agents: in other words, efficient trajectories minimize the total path length, the time needed to complete the task, or the fuel/energy expenditure. In the classical setup, targets locations are known and an assignment strategy is sought that maximizes the global success rate.

During the course of the program, we have considered a class of cooperative motion coordination problems, to which we can refer to as *dynamic vehicle routing*, in which service requests are not known a priori, but are dynamically generated over time by a stochastic process in a geographic region of interest. Specifically, we focus our interest on the Dynamic Traveling Repairperson Problem (DTRP). The m -vehicle DTRP was first studied in [6]. The prototypical problem setup is as follows. Let the environment \mathcal{Q} be a convex, compact set. Consider m identical vehicles which move with bounded speed. Information on outstanding targets at time t is summarized as a finite set of target positions $D(t) \subset \mathcal{Q}$. Targets are generated, and inserted into D , according to a time-invariant spatio-temporal Poisson process, with time intensity λ and some known spatial density. Servicing of a target and its removal from the set D is achieved when one of the vehicles moves to the target location. The objective of the m -DTRP is to minimize the steady-state system time, i.e., the average time that a target has to wait from when it is generated to when it is serviced. Centralized policies were proposed in [6] for the light load ($\lambda \rightarrow 0^+$) and the heavy load case ($\lambda \rightarrow +\infty$) which are within a constant factor of the optimal.

3.1 Dynamic Traveling Repairperson Problem: Decentralized Policies

In [44] we proposed *decentralized* algorithms for these vehicle routing problems. We first designed an optimal policy for the single vehicle, which we call the SINGLE-VEHICLE RECEDING HORIZON MEDIAN/TSP (SRH) POLICY. The policy can be described as follows.

While D is empty, move towards the median of \mathcal{Q} , if not already at it. Let $\text{ETSP}(D)$ be the length of the shortest tour as given by the Euclidean Traveling Salesperson Problem (ETSP) over D . While D is not empty, do the following: (i) for a given $\eta \in (0, 1]$, find a path that maximizes the number of targets reached within $\tau = \max\{\text{diam}(\mathcal{Q}), \text{ETSP}(D)\}$ time units; (ii) service from the current location this optimal fragment. Repeat.

This policy was shown to be asymptotically optimal in the light load case and within a constant factor of the optimal for the heavy load case. The SRH policy was combined with distributed algorithms for locational optimization to give the MULTIPLE-VEHICLE RECEDING HORIZON MEDIAN/TSP (MRH) POLICY. The policy works as follows.

For all $i \in \{1, \dots, m\}$, the i -th vehicle computes its Voronoi cell \mathcal{V}_i and executes the SRH policy in \mathcal{V}_i with the single following modification. While the vehicle is servicing targets in an optimal fragment of $D \cap \mathcal{V}_i$, it will short cut all targets already serviced by other vehicles.

This policy was shown to be locally asymptotically optimal in the light load case, and simulation results suggest that the MRH policy achieves the same performance of the best known centralized policy.

In [3], we proposed control strategies that, while making minimal or no assumptions on communications between vehicles, provide the same level of steady-state performance achieved by the decentralized strategies described above. In other words, we demonstrated that inter-agent communication does not improve the efficiency of such systems, but merely affects the rate of convergence to the steady state. Furthermore, the proposed strategies do not rely on the knowledge of the details of the underlying stochastic process. We also showed that our proposed strategies provide an efficient, pure Nash equilibrium in a game theoretic formulation of the problem, in which each agent's objective is to maximize the number of targets it visits.

3.2 The Dubins Traveling Salesperson Problem (DTSP)

Inspired by many applications of the emergent Unmanned Air Vehicle (UAV) technology, we have investigated a novel class of combinatorial motion planning problems for certain classes of vehicles. One such model is the Dubins model which is a widely accepted kinematic model for fixed-wing aircrafts. A Dubins vehicle is a nonholonomic vehicle that is constrained to move along paths of bounded curvature without reversing direction. We have developed some novel tools and algorithms for optimal motion planning problems for the Dubins vehicle required to visit collections of points in the plane, where the vehicle is said to *visit* a region in the plane if the vehicle goes to that region and passes through it. The objective is to find the shortest path for such vehicle through a given set of target points; we refer to the corresponding problem as the Dubins Traveling Salesperson Problem (DTSP).

The literature on the Dubins vehicle and the TSP is very rich and includes contributions from researchers in multiple disciplines. However, unlike other variations of the TSP, the Dubins TSP cannot be formulated as a problem on a finite-dimensional graph, thus preventing the use of well-established tools in combinatorial optimization.

The DTSP was introduced in our early work [135], where a constant-factor approximation algorithm for the worst-case setting of the DTSP was proposed. In [131], we introduced the stochastic DTSP and gave the first algorithm yielding, with high probability, a solution with a cost upper bounded by a strictly sub-linear function of the number n of target points. Specifically, it was shown that the lower bound on the stochastic DTSP was of order $n^{2/3}$ and that our algorithm performed asymptotically within a $(\log n)^{1/3}$ factor to this lower bound with high probability. In [136]

we designed the first algorithm that asymptotically achieves a constant factor approximation to the stochastic DTSP with high probability. These results were summarized in [137] and [130].

3.2.1 From the Euclidean to the Dubins TSP

Let $\rho > 0$ be the *minimum turning radius* for the Dubins vehicle and let $\text{DTSP}_\rho(P)$ denote the cost of the Dubins TSP over a point set P , i.e., the length of the shortest closed Dubins path through all points in P with minimum turning radius ρ .

One key objective that we addressed was the design of algorithms that provide a provably good approximation to the optimal solution of the Dubins TSP. To establish what a “good approximation” might be, we summarize what is known about the ETSP. The ETSP is of upper and lower order \sqrt{n} as the number of targets n grows, both in the worst case and the stochastic case. Motivated by the Euclidean case, we showed that the DTSP grows with n in the worst case and with $n^{2/3}$ in the stochastic case (as both lower and upper bounds). Most importantly, we proposed novel algorithms for the DTSP in the worst-case and stochastic settings, whose performances are within a constant factor of the optimal solution in the asymptotic limit as $n \rightarrow +\infty$. Finally, we showed the implications of these results in the DTRP for the Dubins vehicle.


3.2.2 The DTSP in the worst case

DTSP lower bound : We first gave a lower bound on $\text{DTSP}_\rho(P)$ in the worst case. For all $\rho > 0$ and $n \geq 2$, we construct in [135] point sets P of n “arbitrarily close” points such that the DTSP in the worst case grows at least linearly in n .

The Alternating Algorithm : Next, we proposed the ALTERNATING ALGORITHM [135] for the DTSP. The underlying principles of the algorithm are the following two observations. First a solution for the DTSP consists of determining the order in which the Dubins vehicle visits the given set of points, and assigning headings for the Dubins vehicle at the points. Second, an approximation for the optimal ordering can be computed by computing an optimal ETSP tour of P . To determine the headings, we use the following alternating heuristics:

- 1: $(a_1, \dots, a_n) :=$ optimal ETSP ordering of P
- 2: $\psi_1 :=$ orientation of segment from a_1 to a_2
- 3: FOR $i \in \{2, \dots, n-1\}$, DO
 - IF i is even, THEN $\psi_i := \psi_{i-1}$, ELSE $\psi_i :=$ orientation of segment from a_i to a_{i+1}
- 4: IF n is even, THEN $\psi_n := \psi_{n-1}$, ELSE $\psi_n :=$ orientation of segment from a_n to a_1
- 5: return the sequence of configurations $\{(a_i, \psi_i)\}_{i \in \{1, \dots, n\}}$.

We illustrate a sample output of this algorithm in the following figures: the left figure is an ETSP solution, the right figure is a Dubins solution generated by the ALTERNATING ALGORITHM.



Analysis of the Alternating Algorithm and DTSP upper bound First, the length of the DTSP in the worst case is upper bounded by the length of ALTERNATING ALGORITHM tour and

hence $\text{DTSP}_\rho(P) \leq \text{ETSP}(P) + \kappa \lceil \frac{n}{2} \rceil \pi \rho$, where $\kappa \simeq 2.6575$. This statement and the lower bound together imply that the DTSP in the worst case grows linearly in n . Additionally, we showed that, in the worst case, the ALTERNATING ALGORITHM performance is within a $\frac{\kappa}{2}$ factor from the optimum as $n \rightarrow +\infty$ and within a $\left(1 + \frac{\kappa\pi}{2\eta}\right)$ factor from the optimum if the minimal inter-target distance is greater than $\eta\rho$, for some $\eta > 0$.


3.2.3 The Stochastic DTSP

Next, we considered a stochastic rather than adversarial placement of the target points. Analogously to the previous discussion, we present a combined design and analysis results (under this new stochasticity assumption). We studied this problem for a rectangular \mathcal{Q} without any loss of generality.

Lower bound In [39], we provided the following lower bound for the stochastic DTSP. For all $\rho > 0$, the expected cost of the DTSP for a set P of n uniformly-randomly-generated points in a rectangle of width W and height H satisfies $\lim_{n \rightarrow +\infty} \frac{\mathbb{E}[\text{DTSP}_\rho(P)]}{n^{2/3}} \geq \frac{3}{4} \sqrt[3]{3\rho WH}$. This bound implies that the stochastic DTSP grows at least with $n^{2/3}$.

A basic geometric construction The key tool in our algorithm design is the following geometric object. Consider two points p_- and p_+ on the plane, with $\ell = \|p_+ - p_-\|_2 \ell_e 4\rho$, and let $\mathcal{B}_\rho(\ell)$ denote the blue region detailed in the figure; we refer to such a region as a *bead* of length ℓ . The region $\mathcal{B}_\rho(\ell)$ enjoys the following properties:

1. its maximum height and its area can be easily computed and are of order ℓ^2 and ℓ^3 as $\ell \rightarrow 0^+$;
2. for any $p \in \mathcal{B}_\rho(\ell)$, there is a Dubins path through the points $\{p_-, p, p_+\}$, entirely contained within $\mathcal{B}_\rho(\ell)$, whose length is at most of order ℓ ,
3. the plane can be periodically tiled by identical copies of $\mathcal{B}_\rho(\ell)$, for $\ell \in]0, 4\rho]$.



The Recursive Bead-Tiling Algorithm The algorithm [136] consists of a sequence of phases;


during each phase, a Dubins tour (i.e., a closed path with bounded curvature) is constructed that “sweeps” the rectangle \mathcal{Q} . We begin by considering a tiling of the plane aligned with the rectangle and such that the area of the bead $\mathcal{B}_\rho(\ell)$ is $WH/(2n)$. In the first phase of the algorithm, a Dubins tour is computed with the following properties:

1. it visits all non-empty beads once,
2. it visits all rows in sequence top-to-down, alternating between left-to-right and right-to-left passes, and visiting all non-empty beads in a row,
3. when visiting a non-empty bead, it services at least one target in it.

In order to visit the targets outstanding after the first phase, a second phase is initiated. Instead of considering single beads, we now consider “meta-beads” composed of two beads each and proceed in a way similar to the first phase, i.e., a Dubins tour is constructed with the following properties:

1. it visits all non-empty meta-beads once,
2. it visits all (meta-bead) rows in sequence top-to-down, alternating between left-to-right and right-to-left passes, and visiting all non-empty meta-beads in a row,
3. when visiting a non-empty meta-bead, it services at least one target in it.

The first, second and third phase are shown in the following figure.



This process is iterated $\lceil \log_2 n \rceil$ times, and at each phase, meta-beads composed of two neighboring meta-beads from the previous phase are considered; in other words, the meta-beads at the i^{th} phase are composed of 2^{i-1} neighboring beads. After the last recursive phase, the leftover targets are visited using the ALTERNATING ALGORITHM in what we call the final phase.

Analysis of the algorithm We first proved a key result which states that the number of outstanding targets after the execution of the $\lceil \log_2 n \rceil$ recursive phases of the RECURSIVE BEAD-TILING ALGORITHM is less than $24 \log_2 n$ with probability one. If $L_{\text{RBTA},\rho}(P)$ denotes the length of the Dubins path computed by the RECURSIVE BEAD-TILING ALGORITHM for a uniformly randomly generated set P of n points in a rectangle of width W and height H , then with probability one $\lim_{n \rightarrow +\infty} \frac{L_{\text{RBTA},\rho}(P)}{n^{2/3}} \leq \ell_e 24 \sqrt[3]{\rho W H} \left(1 + \frac{7}{3} \pi \frac{\rho}{W}\right)$. This statement and the lower bound together imply that the stochastic DTSP grows with $n^{2/3}$. Additionally, in the stochastic setting the RECURSIVE BEAD-TILING ALGORITHM performs within a $(32/\sqrt[3]{3}) \left(1 + \frac{7}{3} \pi \frac{\rho}{W}\right)$ factor from the optimum as $n \rightarrow +\infty$. The computational complexity of the RECURSIVE BEAD-TILING ALGORITHM is of order n and, therefore, the algorithm is easily implementable; in fact, we can compute Dubins tours for thousands of targets in less than a minute.

3.3 The DTRP for Dubins vehicles under heavy load

Here we describe program accomplishments on the Dynamic Traveling Repairperson Problem (DTRP) for the Dubins vehicle. In this section, we consider the case of *heavy load*, i.e., the problem as the time intensity $\lambda \rightarrow +\infty$.

DTRP lower bound We begin with the following lower bound result [39]. For any $\rho > 0$, the system time T_{DTRP} for the DTRP in a rectangle of width W and height H satisfies $\lim_{\lambda \rightarrow +\infty} \frac{T_{\text{DTRP}}}{\lambda} \geq \frac{81}{64} \rho W H$. This result implies that the system time for the Dubins vehicle depends quadratically on the time intensity λ , whereas in the Euclidean case it depends only linearly on it, e.g., see [6].

The Bead Tiling Algorithm The strategy consists of the following steps:

- 1: Tile the plane with beads of length $\ell := \min\{C_{\text{BTA}}/\lambda, 4\rho\}$, $C_{\text{BTA}} = \frac{7-\sqrt{17}}{4} \left(1 + \frac{7}{3} \pi \frac{\rho}{W}\right)^{-1}$
- 2: Traverse all non-empty beads once, visiting one target per bead; Repeat.

We prove that, for any $\rho > 0$ and $\lambda > 0$, the BEAD-TILING ALGORITHM is a stable policy for the DTRP and the resulting system time T_{BTA} satisfies $\lim_{\lambda \rightarrow +\infty} \frac{T_{\text{BTA}}}{\lambda^2} \ell_e 70.5464 \rho W H \left(1 + \frac{7}{3} \pi \frac{\rho}{W}\right)^3$. Hence, the DTRP for Dubins vehicle grows with λ^2 . Additionally, the BEAD-TILING ALGORITHM performs within a constant factor from the optimum as $\lambda \rightarrow +\infty$. While this result confirms that the our algorithm is successful to solve DTRP with stochastically generated targets, here is a **negative** result: there exists no stable policy for the DTRP when the targets are generated in an **adversarial** worst-case fashion with $\lambda \geq (\pi\rho)^{-1}$.

Direct applications to the multi-vehicle DTRP We deal with the multiple-vehicle case in [41] where we suggest the following strategy in the case in which m Dubins vehicles are available. Divide the region into m strips of width $W' = W$ and height $H' = H/m$, and assign one vehicle to each strip. If each vehicle executes the BEAD-TILING ALGORITHM within its own strip, the system time can be computed as $\lim_{\lambda \rightarrow +\infty} \frac{T_{\text{BTA}}}{\lambda^2} \ell_e 70.5464 \frac{\rho W H}{m^3} \left(1 + \frac{7}{3} \pi \frac{\rho}{W}\right)^3$. It is interesting to note that the system time decreases with the *cube* of the number of vehicles. This provides a very **strong** motivation for the use of large-scale groups of mobile vehicles, especially when differential constraints such as bounded curvature play an important role.

3.4 The TSP and DTRP for other vehicle models

The novel tools developed for Dubins vehicle have been extended to solve similar problems for **other** vehicle models. Our work in [134] completed the generalization of the known combinatorial results on the ETSP and DTRP (applicable to systems with single integrator dynamics) to double integrators and Dubins vehicle models. It is interesting to compare our results with the setting where the vehicle is modeled by a single integrator; this setting corresponds to the so-called Euclidean case in combinatorial optimization. In the following table the single integrator results in the first column are taken from [148, 6]; the other results are novel and taken from our work.

	Single integrator	Double integrator	Dubins vehicle
Min. time for TSP tour (worst-case)	$\Theta(n^{1-\frac{1}{d}})$	$\Omega(n^{1-\frac{1}{d}}),$ $O(n^{1-\frac{1}{2d}})$	$\Theta(n)$ ($d = 2, 3$)
Exp. min. time for TSP tour (stochastic)	$\Theta(n^{1-\frac{1}{d}})$	$\Theta(n^{1-\frac{1}{2d-1}})$ w.h.p. ($d = 2, 3$)	$\Theta(n^{1-\frac{1}{2d-1}})$ w.h.p. ($d = 2, 3$)
System time for DTRP (heavy load)	$\Theta(\lambda^{d-1})$ ($d = 2$)	$\Theta(\lambda^{2(d-1)})$ ($d = 2, 3$)	$\Theta(\lambda^{2(d-1)})$ ($d = 2, 3$)

Finally, TSP and DTRP problems for Reeds Shepp cars and differential drive robots were considered in [40].

3.5 Scheduling Multiple Vehicles Dynamically and Bandit Problems

Dynamic multi-vehicle scheduling has also been considered using a related approach, where the goal is to dynamically schedule M vehicles to visit N targets and collect “rewards”. We build on the techniques developed for the multi-armed bandit problem (MABP) and its extension, the restless bandits problem (RBP), to construct scalable policies and efficiently computable performance bounds. Two extensions of these models were considered, which are particularly relevant to the **type** of missions executed by autonomous vehicles: environment with imperfect information, and the addition of switching costs for traveling between targets.

Assume that the N targets are two-state Markov Chains evolving independently according to known and distinct transition probability matrices. When an agent explores site i , it can observe its state without measurement error, and obtain a reward R^i if the site is in the first state. Here, there is no cost for moving the agents between the sites. We want to determine how we should allocate the agents at each time period. It is shown in [65] that the greedy policy which consists in observing the M sites with maximum expected immediate reward is not optimal. In other words, there is a value for the information gained when observing sites which might not be in the first state. We formulated the problem as a particular case of the restless bandits problem, although with partial information, and proposed an alternative index policy following the ideas of Whittle [175]. We gave a closed form expression for the indices of this problem, which moreover can be computed separately for each target. The resulting policy chooses to visit the M sites with greatest indices and can therefore scale extremely well with the size of the problem. Finally, we can efficiently compute a performance bound for problems with up to $N = 3000$ sites and $M = N/20$ vehicles.

The addition of a path planning component, via the introduction of switching costs representing travel distances between the sites, complicates the problem significantly. Indeed, with this modification the nice separable structure of the MABP is lost and one cannot design index policies based on individual sites independently. In [63, 64, 67], we propose a linear programming relaxation for this problem (with perfect information, but more complicated dynamics than in the partial information case above), which can be computed in polynomial time and provides:

- an upper bound on the achievable performance
- an approximation of the reward-to-go which can be used in approximate dynamic programming.

The computation of the one-step lookahead policy using the approximate reward-to-go consists simply in solving at each period a linear assignment problem. The relaxation needs to be computed only once offline. This can be done for 30 sites in about 20 minutes on a standard desktop, independently of the number of vehicles. Experimental results showed a gap of typically less than 15% between the performance bound and the proposed policy.

3.6 Risk-Sensitive Multi-Agent Systems

In recent work the impact of the number of agents on the performance of a system that is risk-sensitive has been considered. In particular, in [68] a tracking problem is considered with a randomly moving evader and n pursuers, which can obtain noisy measurements of their respective separation from the evader. The problem formulation is in the linear exponential quadratic Gaussian (LEQG) framework; i.e., using a risk-sensitive performance measure. We demonstrated through simulation and analysis that the threshold risk parameter above which the performance per agent becomes infinite increases with the number of agents, or in other words, that a minimum number of agents is necessary for the risk sensitive LEQG tracking problem to have a solution. We believe that additional work on the asymptotics and moderate n regime would lead to a better understanding of the impact of the number of agents on the performance of multi vehicle autonomous systems.

4 CONTROL WITH NETWORK GRAPH CONSTRAINTS

The program has had major research accomplishments on the theory of optimal decentralized control, where information passing between agents is specified by a graph structure. Much of conventional controls analysis assumes in contrast that the controllers to be designed all have access to the same measurements. However, with the advent of complex systems, decentralized control has become central, because there are multiple controllers each with access to different information. Examples of

such systems include autonomous automobiles on the freeway, the power distribution grid, spacecraft moving in formation, paper machining, in addition to aerial-vehicle networks.

4.1 Multiplayer Games

One of the pervasive goals of this MURI program has deep exploration of the strong links between traditionally diverse areas such as robustness analysis, protocol design and verification, and cooperative control. Even though these fields have been developed independently by different communities, there are enough conceptual similarities between them to make possible a useful synthesis of the techniques. In all these cases, the main conceptual objective is to guarantee that a clearly defined set of “bad behaviors” is avoided. For example, in the case of robustness analysis of linear systems, that set can correspond to a particular combination of uncertain parameters producing an unstable closed-loop behavior, where the signal values diverge to infinity. In protocol verification the bad behavior can be associated, for instance, to a deadlock condition. In motion coordination, this could correspond to a multi-robot collision.

We have been, and continue to be, particularly interested in adversarial situations, where there are several decision makers with possibly conflicting objectives. These situations can be profitably analyzed within the framework of *game theory*. The objective here is to characterize the optimal strategies of the decision makers. Game theory subsumes many aspects of optimization, since that situation corresponds to the case of a single decision maker. We are mainly interested in classes of games where the decision makers have an infinite number of pure strategies to choose from (they can also randomize over this choice). Typical examples of these situations are pursuit-evasion games, and resource allocation in networking. In many games, optimality of the strategies may be too difficult a goal to achieve, and instead we may want to settle for solutions that are “good enough.”

Generally, the effective certification of this kind of properties (optimality, safety, robustness, etc) is very problem dependent. In the cases where the underlying constraints and dynamics of the system are described using polynomials, this opens up the door to using algebraic methods for the efficient verification and certification. The exciting part is that the search for *short proof certificates* can be carried out in an *algorithmic way*. This is achieved by coupling efficient optimization methods and powerful theorems in semialgebraic geometry. For practical reasons, we are only be interested in the cases where we can find *short proofs*, i.e., those that can be verified in polynomial time. A priori, there are no guarantees that a given problem has a short proof. However, in general we can find short proofs that provide useful information: for instance, in the case of optimization problems, this procedure provides lower bounds on the value of the optimal solution.

In the case of polynomials, the central piece of the puzzle is the key role played by *sums of squares decompositions*. The principal numerical tool used in the search for certificates is *semidefinite programming*, a broad generalization of linear and convex quadratic optimization. Semidefinite programs, also known as Linear Matrix Inequalities (LMI) methods, are convex optimization problems, and correspond to the particular case of the convex set being the intersection of an affine family of matrices and the positive semidefinite cone. It is well known that semidefinite programs can be efficiently solved both theoretically and in practice.

Building upon the powerful methods from convex optimization (in particular, sum of squares and semidefinite programming), we have provided novel, effective and efficient solutions to a wide variety of continuous games. Starting with our initial work on minimax equilibria for two-player zero-sum games, we have also analyzed stochastic games, as well as Aumann’s celebrated notion of correlated equilibria for the case of multiplayer games.

4.1.1 Polynomial and semialgebraic games

As part of this MURI, the following conference publications [107, 149, 144, 102] have been completed. Journal versions of several of these have been submitted, or are currently under preparation. We

discuss them in more detail below.

- Polynomial games [107]: A ubiquitous mathematical model of adversarial situations is given by two-person zero-sum games. Usually, these are modeled through finite bimatrix games, where each player has access to finitely many pure strategies. However, in many situations (e.g., pursuit-evasion games, power/rate allocation) it is often the case that there is a continuum of possible strategies for each player. In this work, we initiated the study of zero-sum games where the payoff function is a polynomial expression in the actions of the players. This class of games (“polynomial games”) was originally introduced by Dresher, Karlin, and Shapley in 1950, in their pioneering work in the RAND corporation. We have shown that the value of these games, and the corresponding optimal strategies, can be obtained by solving a single semidefinite programming problem. In addition, we have shown how the results extend, with suitable modifications, to a general class of semialgebraic games.
- Separable games [149]: these are a structured subclass of continuous games, whose payoffs take a sum-of-products form. This subclass includes all finite games and polynomial games. Separable games provide a unified framework for analyzing and generating results about the structural properties of low rank games. This work extends previous results on low-rank finite games by allowing for multiple players and a broader class of payoff functions. We have introduced methods for the computation of equilibria, and connected these results with alternative characterizations of separability that show that separable games are the largest class of continuous games to which low-rank arguments apply.
- Stochastic games [144]: Stochastic games simultaneously generalize “standard” multiplayer games, and Markov decision processes. In a stochastic game, the players’ actions affect not only their immediate payoff, but also the transition probabilities that define the next state of the game. Thus, players should carefully balance their immediate rewards, versus the long-term objective of remaining in a favorable situation. In this work, we consider finite state two-player zero-sum stochastic games over an infinite time horizon with discounted rewards. As in our earlier work, the players have infinite strategy spaces and the payoffs are assumed to be polynomials. In order to obtain tractable results, we have restricted our attention to a special class of games for which the “single-controller” assumption holds. This assumption implies that only one of the players directly affects the transition probabilities. Our main result in this paper is a characterization of the minimax equilibria and optimal strategies via SOS and semidefinite programming.
- Correlated equilibria in multiplayer games [102]: The classical equilibrium notion for multiplayer games is that of Nash equilibria. More recently, Aumann’s alternative definition of *correlated equilibrium* has received much attention as a generalization of the Nash solution, which is both justifiable in theory and efficiently computable in practice. The idea of a correlated equilibrium is that each player receives a private recommendation of what strategy to play, but these recommendations may be correlated. If all the players know the joint distribution of the recommendations, then they can each compute the joint conditional distribution of their opponent’s recommendations given their own recommendation. If each player’s recommendation is always a best response to this conditional distribution, then the distribution of recommendations is called a correlated equilibrium (the Nash solution is recovered if additionally the recommendations to each player are independent). In this work, we consider the problems of characterizing and computing correlated equilibria in polynomial games with infinite strategy sets. We prove several characterizations of correlated equilibria in continuous games which are more analytically tractable than the standard definition. Then we use these to construct algorithms for approximating correlated equilibria of polynomial games with arbitrary accuracy, including a sequence of semidefinite programming relaxation algorithms and discretization algorithms.

4.2 Decentralized Control

The paper [125] addresses the problem of constructing optimal decentralized controllers. The problem is formulated as one of minimizing the closed-loop norm of a feedback system subject to constraints on the controller structure. The paper defines the important notion of quadratic invariance of a constraint set with respect to a system, and show that if the constraint set has this property, then the constrained minimum-norm problem may be solved via convex programming. It also shows that quadratic invariance is necessary and sufficient for the constraint set to be preserved under feedback. These results are developed in a very general framework, and are shown to hold in both continuous and discrete time, for both stable and unstable systems, and for any norm. This notion unifies many previous results identifying specific tractable decentralized control problems, and delineates the largest known class of convex problems in decentralized control.

A number of useful and practical examples of this theory have been produced. For example, optimal stabilizing controllers may be efficiently computed in the case where distributed controllers can communicate faster than their dynamics propagate. This research has provided a test for sparsity constraints to be quadratically invariant, and thus amenable to convex synthesis.

In a standard controls framework, the decentralization of the system manifests itself as sparsity or delay constraints on the controller to be designed. Therefore a canonical problem one would like to solve in decentralized control is to minimize a norm of the closed-loop map subject to a subspace constraint as follows

$$\begin{aligned} & \text{minimize} && \|f(P, K)\| \\ & \text{subject to} && K \text{ stabilizes } P \\ & && K \in S \end{aligned}$$

Here $\|\cdot\|$ is any norm on $\mathcal{R}_p^{n_z \times n_w}$, chosen to encapsulate the control performance objectives, and S is a subspace of admissible controllers which encapsulates the decentralized nature of the system. The norm on may be either a deterministic measure of performance, such as the induced norm, or a stochastic measure of performance, such as the \mathcal{H}_2 norm. Many decentralized control problems may be formulated in this form, and some examples are shown below. The subspace S is called the *information constraint*.

This problem is made substantially more difficult in general by the constraint that K lie in the subspace S . Without this constraint, the problem may be solved by a simple change of variables. For specific norms, the problem may also be solved using a state-space approach. Note that the cost function $\|f(P, K)\|$ is in general a non-convex function of K .

For a general linear time-invariant plant P and subspace S there is no known tractable algorithm for computing the optimal K . It has been known since 1968 that even the simplest versions of this problem can be extremely difficult. In fact, certain cases have been shown to be intractable. However, there are also several special cases of this problem for which efficient algorithms have been found. The paper unifies these cases and identifies a simple condition, called *quadratic invariance*, under which the above problem may be recast as a convex optimization problem. The notion of quadratic invariance allows one to better understand this dichotomy between tractable and intractable optimal decentralized control problems. It further delineates the largest known class of decentralized problems for which optimal controllers may be efficiently synthesized.

Quadratic invariance is a simple algebraic condition relating the plant and the constraint set. The main results of [125] hold for continuous-time or discrete-time systems, for stable or unstable plants, and for the minimization of any norm.

It is also worth noting that that optimal synthesis of a symmetric controller for a symmetric plant is also quadratically invariant and thus amenable to convex synthesis. This is important because this problem, while formerly known to be solvable, defied other efforts to classify tractable problems.

The paper [125] develops an explicit test for the quadratic invariance of sparsity constraints, and thus shows that optimal synthesis subject to such constraints which pass the test may be cast as a

convex optimization problem. A consequence of the test is that block diagonal constraints are **never** quadratically invariant unless the plant is block diagonal as well.

These results all hold for the minimization of an arbitrary norm. If the norm of interest is the \mathcal{H}_2 -norm, then the constrained convex optimization problem may be further reduced to an unconstrained convex optimization problem, and then readily solved.

Quadratic invariance. The characterization of constraint sets S that lead to tractable solutions for the decentralized control problem is as follows.

Definition 3. Suppose $G \in \mathcal{L}(\mathcal{U}, \mathcal{Y})$, and $S \subseteq \mathcal{L}(\mathcal{Y}, \mathcal{U})$. The set S is called quadratically invariant under G if

$$KGK \in S \quad \text{for all } K \in S$$

Roughly speaking, if the set S is quadratically invariant, then optimal control synthesis subject to the constraint that K lie in S may be solved via convex program.

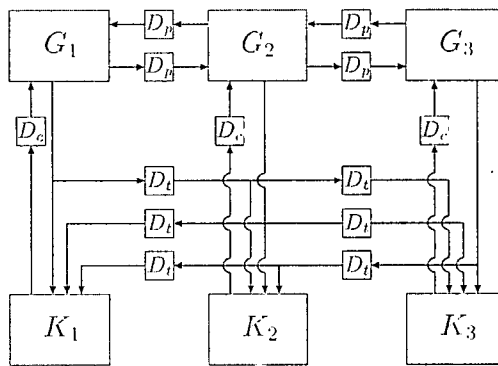


Figure 7: Distributed Control Problem

Distributed control with delays. One particular distributed control problem is shown in Figure 7. Suppose there are n subsystems with transmission delay $t \geq 0$, propagation delay $p \geq 0$ and computational delay $c \geq 0$. If

$$t \leq p + \frac{c}{(n-1)}$$

then the corresponding set S is quadratically invariant under the corresponding G . Hence finding the minimum-norm controller may be reduced to a convex optimization problem when the controllers can transmit information faster than the dynamics propagate; that is, when $t \leq p$. One also sees that the presence of computational delay causes this condition to be surprisingly relaxed. This result has been generalized considerably [118].

4.3 Distributed Control of Heterogenous Systems

In [38] we developed an alternative approach to topology constrained control synthesis, over grid lattices, using a generalization of the multidimensional Roesser model. With the assumption that controllers have the same interconnection structure as the nominal system we were able to show that a relaxation of the global optimization could be explicitly reduced to a semidefinite program. More specifically, we extended optimal control machinery to include heterogeneous Roesser systems, and derive sufficient conditions for analyzing performance with respect to the induced 2-norm, and provided sufficient conditions for the existence controllers which stabilize the system and provide

a guaranteed level of performance (this is the same global performance criterion introduced in the preceding section). The techniques developed are based on extending and combining those developed by the PIs and co-workers on nonstationary systems and homogeneous distributed systems.

In recent work [37] we consider arbitrary graphs, and show how the approach in [38], which is restricted to grid topologies, can be extended to address general interconnection topologies. The new framework is able to capture arbitrary graphs and provides a unifying view and approach to all previous work on using Roeser-type models for distributed control.

4.4 Decentralized Control of Markov Processes

Decentralized decision problems are optimization problems in which a collection of decisions are made in response to a set of observations with the goal of minimizing some cost. The complicating factor is that each decision is made based only on knowledge of a subset of the observations. That is, the complete set of observations can be thought of as the state of the environment. Each decision is made on the basis of an incomplete observation of the state, although the cost incurred depends on the entire state and set of decisions. Such problems are common in areas such as engineering and economics. Much of the early work on decentralized decision problems was motivated by economic problems. In certain engineering problems, such as the design of distributed detection schemes and distributed data transmission protocols, the key difficulty lies in the design of good rules for interacting decision makers to follow.

The paper [26] considers fairly general discrete versions of this problem, where the sets of possible observations and decisions are finite. The first problem considered is a *static* decision problem, where a single set of decisions is made in response to a single set of observations. Given the probabilities of all sets of observations, the goal is to choose decentralized decision rules which minimize the expected cost. This problem is known to be \mathcal{NP} -hard, even for the case of two decision makers. Therefore, our goal is to determine effective methods of computing good *suboptimal* solutions to this problem. The paper shows that this problem can be equivalently formulated as a minimization of a polynomial subject to linear constraints. Relaxations of this polynomial optimization problem can then be efficiently solved. From these relaxations, one obtains lower bounds on the minimum achievable value for the original problem, as well as suboptimal decision rules. The combination of lower bounds together with suboptimal solutions is powerful, since this gives us a way to put a bound on how suboptimal the best known decision rules are.

The paper also extends that analysis to a *dynamic* version of the decentralized decision problem. In this problem, a sequence of observations and decisions are made, and the sequence of observations evolves according to a Markov chain determined by the decisions. The goal in such a problem is to choose decentralized decision rules to minimize the average steady-state cost earned by the system. The complexity results for the static problem are extended to show that this problem is \mathcal{NP} -hard as well. As for the static problem, one can formulate this problem as a polynomial optimization problem and solve relaxations to obtain lower bounds on the minimum cost for the original problem as well as suboptimal decision rules.

The problem of *decentralized detection* is an example of a decentralized stochastic decision problem. In a detection problem, there are several hypotheses on the underlying state of our environment, and one would like use measurements of our environment to decide which hypothesis is true.

Classical detection methods assume all measurements are available to a single detector, which estimates the true hypothesis based on all measurements. Such a detection scheme is called *centralized*. Optimal decision rules in centralized schemes are given by the well-known MAP (maximum a-posteriori probability) detector. In a *decentralized* detection scheme, each sensor is responsible for making a decision based only on its own measurement. The goal is to choose decision rules for all sensors which are optimal with respect to some system-wide cost function.

For example, suppose one has a collection of sensors each monitoring various elements of some industrial process. One would like the sensors to sound an alarm when some part of the process is




Figure 8: The correct hypothesis $H \in \{h_1, \dots, h_M\}$ is to be detected. In this figure, N independent detectors produce decisions u_i based on their measurements y_i .

malfunctioning. In this case one may wish to maximize the probability that the alarm sounds **when** there is a malfunction and does not sound when there is no malfunction. One option is to **transmit all** sensor measurements to a central location, where a decision to sound the alarm is made **on the basis** of all measurements. An alternative is to equip each sensor with its own decision rule and **the ability** to sound the alarm. When the loss of performance associated with employing the second **alternative** is small, such a scheme is preferable due to the reduced implementation complexity associated **with** the elimination of the communication requirements.

One might initially assume that good decentralized decision rules can be obtained **by** allowing each sensor to use a MAP detection rule. While this is true in some special cases, it is **not true** in general. Unlike the centralized case, the general problem of computing optimal decentralized detection rules is \mathcal{NP} -hard. Also, decentralized decision rules can appear considerably more complex than their centralized counterparts. For example, optimal decentralized decision rules typically involve hedging among the sensors, a strategic element which is not present when simply using MAP rules at each detector.

Due to the complexity of this problem, most existing methods for computing decentralized detection rules produce locally optimal equilibrium policies. Such policies are said to be *person-by-person optimal*: for a set of such decision rules, no improvement can be obtained by adjusting the decision rule for any given sensor while leaving the others fixed. In general, a single problem instance may have many equilibrium policies. The globally optimal policy is clearly an equilibrium policy. However, for any given equilibrium policy, one has no way of knowing how this policy relates to the globally optimal policy. In particular, one has no way of knowing how much improvement could be obtained by using the globally optimal policy.

The paper [26] shows that an equilibrium policy can perform arbitrarily poorly compared to the optimal policy. The methods developed in that paper are *relaxations*. In addition to generating an equilibrium policy, they return a lower bound on the minimum achievable cost by any decentralized policy. When the bound is exact, one has a proof that our computed policy is globally optimal. Even when the bound is not exact, one has a measure of the suboptimality of the computed policy. In [71] we study a special decentralized detection problem which surprisingly yields a **thresh-hold-type** policy.

5 VERIFICATION AND VALIDATION FOR COMPLEX SYSTEMS

To realize the control-level algorithms and strategies of the preceding sections in real multi-vehicle systems, they need to be parlayed, through engineering, into detailed hardware and software implementations. The accomplishments described in the current section — models, tools and techniques — are aimed at verifying and validating such complex engineered implementations. One of the key goals of this project was to identify and develop such techniques so as to formally verify networked systems for correctness, as a means to reliable design. (The results also apply to verification and validation of implementations related to work developed in Section 6.)

5.1 Input-Output Automata Modeling Frameworks and Methodology

One of the major approaches for verification and validation that was developed in the project was that of input-output automata (IOA). The starting point for work on input-output automata was the foundational work on a modeling framework for hybrid (continuous/discrete) systems, which we call the Hybrid I/O Automata (HIOA) framework [84]. This foundational work was continued with a monograph collecting and summarizing prior work on Timed I/O Automata, and with several papers on Probabilistic I/O Automata and on combined Probabilistic/Timed I/O Automata. Together, this work provides a foundation for modeling and analyzing a wide range of systems, including wired and wireless communication networks, and controlled vehicles and robots.

5.1.1 Mathematical foundations

Hybrid and timed IOA models In [84] the Hybrid Input/Output Automaton (HIOA) modeling framework is presented providing a basic mathematical framework to support description and analysis of hybrid systems. An important feature of this model is its support for decomposing hybrid system descriptions. In particular, the framework includes a notion of external behavior for a hybrid I/O automaton, which captures its discrete and continuous interactions with its environment. The framework also defines what it means for one HIOA to implement another, based on an inclusion relationship between their external behavior sets, and defines a notion of simulation, which provides a sufficient condition for demonstrating implementation relationships. The framework also includes a composition operation for HIOAs, which respects external behavior, and a notion of receptiveness, which implies that an HIOA does not block the passage of time. *The framework is intended to support analysis methods from both computer science and control theory.*

In the monograph [58] the Timed Input/Output Automaton (TIOA) modeling framework is developed, a basic mathematical framework to support description and analysis of timed (computing) systems. Timed systems are systems in which desirable correctness or performance properties of the system depend on the timing of events, not just on the order of their occurrence. Timed systems are employed in a wide range of domains including communications, embedded systems, real-time operating systems, and automated control. Many applications involving timed systems have strong safety, reliability and predictability requirements, which makes it important to have methods for systematic design of systems and rigorous analysis of timing-dependent behavior. An important feature of the TIOA framework is its support for decomposing timed system descriptions. In particular, the framework includes a notion of external behavior for a timed I/O automaton, which captures its discrete interactions with its environment.

Probabilistic IOA models Probabilistic automata (PAs) constitute a general framework for modeling and analyzing discrete event systems that exhibit both nondeterministic and probabilistic behavior, such as distributed algorithms and network protocols; an example are the dynamic resource allocation algorithms discussed above. The behavior of PAs is commonly defined using

schedulers (also called adversaries or strategies), which resolve all nondeterministic choices based on past history. From the resulting purely probabilistic structures, trace distributions can be extracted, whose intent is to capture the observable behavior of a PA. However, when PAs are composed via an (asynchronous) parallel composition operator, a global scheduler may establish strong correlations between the behavior of system components and, for example, resolve nondeterministic choices in one PA based on the outcome of probabilistic choices in the other. It was well known that, as a result of this, the (linear-time) trace distribution precongruence is not compositional for PAs, and that (branchingtime) probabilistic simulation preorder is compositional for PAs. In [86] we establish that the simulation preorder is in fact the coarsest refinement of the trace distribution preorder that is compositional. In [85] we establish that on the domain of probabilistic automata, the trace distribution preorder coincides with the simulation preorder.

We have also studied switched probabilistic input/output automata (or switched PIOA), augmenting the original PIOA framework with an explicit control exchange mechanism [21]. Using this mechanism, we model a network of processes passing a single token among them, so that the location of this token determines which process is scheduled to make the next move. This token structure therefore implements a distributed scheduling scheme: scheduling decisions are always made by the (unique) active component. Distributed scheduling allows us to draw a clear line between local and global nondeterministic choices. We then require that local nondeterministic choices are resolved using strictly local information. This eliminates unrealistic schedules that arise under the more common centralized scheduling scheme. As a result, we are able to prove that our trace-style semantics is compositional. We also propose switch extensions of an arbitrary PIOA and use these extensions to define a new trace-based semantics for PIOAs [20].

We introduce the notion of approximate implementations for Probabilistic I/O Automata (PIOA) in [96], and develop methods for proving such relationships. We employ a task structure on the locally controlled actions and a task scheduler to resolve nondeterminism. The interaction between a scheduler and an automaton gives rise to a trace distributional probability distribution over the set of traces. We define a PIOA to be a (discounted) approximate implementation of another PIOA if the set of trace distributions produced by the first is close to that of the latter, where closeness is measured by the (resp. discounted) uniform metric over trace distributions. We propose simulation functions for proving approximate implementations corresponding to each of the above types of approximate implementation relations. Since our notion of similarity of traces is based on a metric on trace distributions, we do not require the state spaces nor the space of external actions of the automata to be metric spaces.

A Probabilistic I/O Automaton (PIOA) is a countable-state automaton model that allows nondeterministic and probabilistic choices in state transitions. A task-PIOA adds a task structure on the locally controlled actions of a PIOA as a means for restricting the nondeterminism in the model. The task-PIOA framework defines exact implementation relations based on inclusion of sets of trace distributions. In [95] we develop the theory of approximate implementations and equivalences for task-PIOAs. We propose a new kind of approximate simulation between task-PIOAs and prove that it is sound with respect to approximate implementations. Our notion of similarity of traces is based on a metric on trace distributions and therefore, does not require the state spaces nor the space of external actions (output alphabet) of the underlying automata to be metric spaces. This work has direct application to approximate implementations to probabilistic safety verification.

Probabilistic and timed IOA models Probabilistic Timed I/O Automaton (PTIOA) framework for modelling and analyzing discretely communicating probabilistic hybrid systems is developed in [97]. State transition of a PTIOA can be nondeterministic or probabilistic. Probabilistic choices can be based on continuous distributions. Continuous evolution of a PTIOA is purely nondeterministic. PTIOAs can communicate through shared actions. By supporting external nondeterminism, the framework allows us to model arbitrary interleaving of concurrently executing automata. The framework generalizes several previously studied automata models of its class. We developed trace-based

semantics for PTIOAs which involves measure theoretic constructions on the space of executions of the automata. We introduce a new notion of external behavior for PTIOAs and show that PTIOAs have simple compositionality properties with respect this external behavior.

5.1.2 Computer-assisted tools and formal techniques

Our work on basic modeling frameworks has been supported by work on formal analysis methods and computer-assisted tools. These include the preliminary modeling language IOA, and the newer, more professional language TIOA (see www.veromodo.com). TIOA tools include the language and front end, a simulator, and translators to the PVS theorem-prover, and the UPPAAL model-checker. Some of this work was also supported by the AFOSR under two STTR contracts. We also include in this category some work involving the development of strategies for using PVS to analyze systems.

IOA toolkit. The IOA Toolkit supports a range of validation methods, including simulation and machine-checked proofs. The manual [53] and reference guide defines the IOA language. Part I of the thesis [154] presents a strategy for compiling distributed systems specified in IOA into Java programs running on a group of networked workstations. IOA is a formal language for describing distributed systems as I/O automata. The translation works node-by-node, translating IOA programs into Java classes that communicate using the Message Passing Interface (MPI). The resulting system runs without any global synchronization. We proved that, subject to certain restrictions on the program to be compiled, assumptions on the correctness of hand-coded datatype implementations, and basic assumptions about the behavior of the network, the compilation method preserves safety properties of the IOA program in the generated Java code. We model the generated Java code itself as a threaded, low-level I/O automaton and use a refinement mapping to show that the external behavior of the system is preserved by the translation. The IOA compiler has been implemented at MIT as part of the IOA toolkit. The toolkit supports algorithm design, development, testing, and formal verification using automated tools. The IOA language provides notations for defining both primitive and composite I/O automata. Part II of this thesis describes, both formally and with examples, the constraints on these definitions, the composability requirements for the components of a composite automaton, and the transformation a definition of a composite automaton into a definition of an equivalent primitive automaton.

The capabilities and performance of the IOA Toolkit are reported in [54], and in particular the tools that provide support for implementing and running distributed systems (checker, composer, code generator). The Toolkit compiles distributed systems specified in IOA into Java classes, which run on a network of workstations and communicate using the Message Passing Interface (MPI). In order to test the toolkit, several distributed algorithms were implemented, ranging from simple algorithms such as LCR leader election in a ring network to more complex algorithms such as the GHS algorithm for computing the minimum spanning tree in an arbitrary graph.

A strategy for compiling distributed systems specified in IOA is summarized in [156], a formal language for describing such systems as I/O automata, into Java programs running on a group of networked workstations. The translation works node-by-node, translating IOA programs into Java classes that communicate using the Message Passing Interface. The resulting system runs without any global synchronization. We prove that, subject to certain restrictions on the program to be compiled, assumptions on the correctness of hand-coded datatype implementations, and basic assumptions about the behavior of the network, the compilation method preserves safety properties of the IOA program in the generated Java code. We model the generated Java code itself as a threaded, low-level I/O automaton and use a refinement mapping to show that the external behavior of the system is preserved by the translation. The IOA compiler is part of the IOA toolkit which supports algorithm design, development, testing, and formal verification using automated tools.

The IOA language developed and reported in [155] provides notations for defining both primitive and composite I/O automata. This note describes, both formally and with examples, the constraints

on these definitions, the composability requirements for the components of a composite automaton, and the transformation of a composite automaton into an equivalent primitive automaton.

In [167] we describe our approach and design for code generation that focuses on the issue of removing implicit nondeterminism and specify a transformation on IOA programs that makes all nondeterminism explicit. The programmer can then replace all explicit nondeterminism with deterministic statements prior to code generation. We also describe this transformation at a semantic level i.e., at the level of the I/O automaton mathematical model. It is shown that the transformation defined at the IOA level conforms to the one at the semantic level.

The thesis [147] concerns the addition of a capability to simulate composite automata in a manner that allows observing and debugging the individual system component automata. While there is work in progress on creating a tool that will translate a composite definition into a single automaton, the added ability to simulate composite automata directly will add modularity and simplicity, as well as ease of observing the behavior of individual components for the purpose of distributed debugging.

PVS strategies. A related support tool based on the PVS theorem prover that can help users prove a candidate abstraction relation correct. This tool support relies on a clean and uniform technique for dening abstraction properties relating automata that uses library theories for defining abstraction relations and templates for specifying automata and abstraction theorems. The work is reported in [92] describes how the templates and theories allow development of generic, high level PVS strategies that aid in the mechanization of abstraction proofs. These strategies first set up the standard subgoals for the abstraction proofs and then execute the standard initial proof steps for these subgoals, thus making the process of proving abstraction properties in PVS more automated. Two types of abstraction properties are the focus: refinement and forward simulation.

We have also developed an abstraction specification technique and associated abstraction proof strategies we are developing for I/O automata [91]. The new strategies can be used together with existing strategies in the TAME (Timed Automata Modeling Environment) interface to PVS; thus, our new templates and strategies provide an extension to TAME for proofs of abstraction. We have extended the set of TAME templates and strategies.

The toolkit has been expanded to handle TIOA models, and is aimed at supporting system development based on TIOA specifications [2]. The TIOA toolkit is an extension of the IOA toolkit, which provides a specification simulator, a code generator, and both model checking and theorem proving support for analyzing specifications. Also, several illustrative examples are provided in [2].

5.1.3 Vehicle-based examples and applications

As part of the project, case studies were performed drawing example systems from the general area of automated vehicles and robot control, and modeled and analyzed using Timed and Hybrid I/O Automata theory described above. As part of this effort a type of abstraction layer for programming mobile networks, which we call "Virtual Node layers", was developed. Virtual Node layers can be used for coordination applications; we have explored three kinds, namely, robot pattern formation, a simple intelligent highway application (a Virtual Traffic light), and a simple air-traffic-control system (based on Virtual ATCs).

Quanser helicopter case study In [99] a case study of the developed hybrid verification is provided, and a formal verification of the safety properties of NASA's Small Aircraft Transportation System (SATS) landing protocol is carried out. A new model is presented using the timed I/O automata (TIOA) framework [58], and key safety properties are verified. Properties specific to the new model, such as lower bounds on the spacing of aircraft in specific areas of the airspace, are provided.

The Hybrid I/O Automaton modelling framework [84] is applied to a realistic hybrid system verification problem in [98]. A supervisory pitch controller for ensuring the safety of a model helicopter

system is designed and verified. The supervisor periodically observes the plant state and takes over control from the user when the latter is capable of taking the plant to an unsafe state. The paper also presents a set of language constructs for specifying hybrid I/O automata.

Tracking Stalk, a hierarchy-based fault-local stabilizing algorithm for tracking in sensor networks, is developed in [36]. Starting from an arbitrarily corrupted state, Stalk satisfies its specification within time and communication cost proportional to the size of the faulty region instead of the network size. Local stabilization is achieved by slowing propagation of information as the levels of the hierarchy underlying Stalk increase, enabling the more recent information propagated by lower levels to override misinformation at higher levels. While achieving fault-local stabilization, Stalk also adheres to the locality of tracking operations: an operation to find a mobile object at a distance d away requires $O(d)$ amount of time and communication cost to intercept the moving object, and a move of an object to a distance d away requires $O(d * \log(\text{network diameter}))$ amount of time and communication cost to update the tracking structure. Furthermore, Stalk achieves seamless tracking of a continuously moving object by enabling concurrent executions of move and find operations.

Air-traffic control An assertional-style verification of the aircraft landing protocol of NASA's SATS (Small Aircraft Transportation System) concept using the I/O automata framework and the PVS theorem prover was developed in [160]. An IOA model of the landing protocol was developed, and translated into a corresponding PVS specification; a verification of the safety properties of the protocol using the assertional proof technique and the PVS theorem prover was then successfully performed. A more extensive account can be found in [158].

Virtual node based coordination algorithms In [83] a virtual node abstraction layer was used to coordinate the motion of real mobile nodes in a region of 2-space. In particular, how nodes in a mobile ad hoc network can arrange themselves along a predetermined closed curve in the plane, and can maintain themselves in such a configuration in the presence of changes in the underlying mobile ad hoc network, was considered. The strategy illustrated was allowing the mobile nodes to implement a virtual layer consisting of mobile client nodes, stationary virtual nodes (VNs) at predetermined locations in the plane, and local broadcast communication. The VNs coordinate among themselves to distribute the client nodes relatively evenly among the VNs' regions, and each VN directs its local client nodes to form themselves into the local portion of the target curve.

A general VNLayer architecture was introduced in [8], and then used to design a practical VN-Layer implementation, optimized for real-world use. Discussed [8] also is experience with deploying this implementation on a testbed of hand-held computers, and in a custom-built packet-level simulator, and present a sample application - a virtual traffic light - to highlight the power and utility of our abstraction. The idea of using Virtual Stationary Automata (VSAs) to take a distributed approach to automated air traffic control was extensively explored in [7].

Security protocol modeling and analysis This MURI also brought together team members to apply the techniques developed to security issues. This resulted in a full model and analysis for the Goldreich et al. Oblivious Transfer protocol. This required us to develop a new approach to handling the combination of nondeterministic and probabilistic choice, where nondeterministic choices are resolved independently of probabilistic choices. (If nondeterministic choices are allowed to depend on the results of probabilistic choices, secrets can be divulged unintentionally.) This approach is embodied in our new Task-PIOA modeling framework—an extension of the previous PIOA framework. We extended simulation relation methods for PIOAs to Task-PIOAs. Furthermore, we extended Task-PIOAs so they can express computational limitations, such as polynomial time bounds. In further work, we formulated Canetti's notion of "secure emulation" within the Task-PIOA framework, and proved suitable protocol composition theorems. Currently, we are working

on extending the framework still further, to permit us to analyze “long-lived” security protocols [16, 14, 15, 13, 17, 90]

5.2 Learning and Randomization

We have also approached verification by using techniques of *model checking*. Computer-aided verification is concerned with determining whether a formal model of a system, often presented as a graph of states and state transitions, satisfies certain correctness properties. The most popular algorithmic technique, model checking [22], works by systematically stepping through the global states of the system while checking various properties at each stage. The widespread use of model checking in practice is predicated on two observations: first, the technique is largely automated, and requires limited user input; second, when a system is found to not meet its correctness requirements, the tool provides a counter-example witnessing this, which has been found to be very beneficial in fixing the flaw.

Applying model checking to peer-to-peer networks of vehicle systems operating in a unknown, potentially malicious environment, presents unique challenges. Formal models of such systems have certain fundamental aspects that must be considered. First, there are geographically disperse parties that concurrently compute and communicate. Second, the uncertain environment requires modelling probabilistic events and stochastic behavior. Finally, the individual embedded systems have both discrete and continuous dynamics that have a non-trivial dependence on real-time. Thus, the semantics of such dynamic multi-vehicular systems must be described by transition systems that have *infinitely* many global states — the multitude of global states arise from the need to model the (potentially) unbounded number of messages that have been sent but that have as yet not been delivered, to model real-time and clocks, and to model the many real-valued continuous variables like position, speed and acceleration that are critically used in describing the dynamics.

However, results from theoretical computer science demonstrate that the verification of even extremely simple properties of such infinite state systems is *undecidable*. In other words, there is no mechanized procedure that can automatically verify such systems. In the face of such extremely pessimistic news, researchers have taken two approaches. First, they have come up with *semi-decision* procedures that are not guaranteed to terminate on all systems, but for systems on which they do terminate, are known to give correct and useful answers. Based on such semi-decision procedures, analysis tools have been built and have been used to formally analyze many of the systems that arise in practice. Second, special features present in systems of interest have been identified that manifest themselves in unique structural properties in the state-space of the system, which can then be exploited to yield *decision procedures* for such systems. These algorithms (that always terminate with the right answer) have then been used to formally verify such special systems.

As part of this project, we have pursued both these general research themes to address the unique challenges posed by networks of vehicle systems. We have developed semi-decision procedures for analyzing distributed, probabilistic systems based on two novel paradigms: *learning*, and *randomization*. Next, we have identified general classes of distributed and hybrid systems that have special features that have decidable verification problems. In what follows, we give more details about these accomplishments.

5.2.1 Learning to verify

Symbolic model checking [22], is an algorithmic technique for verification that has been extremely useful in practice. The main thesis of this approach is to observe that verification can be viewed as computing the fixed point of a function. For example, if we want to check if an invariant holds in a system, then the verification problem involves computing the set of *reachable states* (states of the system encountered during some computation), which is a fixed point of the one-step *transition relation*, and checking if all reachable state satisfy the invariant. Now if verification is nothing but

fixed point computation, then one simple algorithm, based on Tarski’s method, is to repeatedly apply the function whose fixed point one is computing, until this process stabilizes. The next key observation in symbolic model checking is to use *symbolic representation* of sets (instead of an explicit representation) during these fixed point computations. The widespread success of this approach is predicated on the observation that practical systems typically have extremely structured fixed points, and thus have very small symbolic representations.

As part of this project, we have initiated and pursued a method that is a radical departure from this traditional approach. Instead of using Tarski’s iterative approach for computing the symbolic representation of the fixed point, in *learning to verify*, we view the model checking problem as a learning problem, and try to *learn* the symbolic representation of the fixed point by observing executions of the system.

This learning based approach has many theoretical and practical benefits. First, it can be used to identify new classes of infinite state systems for which the model checking problem is provably decidable. Second, the running time of the algorithm only depends on the size of the symbolic representation of the fixed point. This is significant because the algorithm can therefore be applied to infinite state system — the fixed point sets for infinite state systems, even though of infinite cardinality, often have a finite symbolic representation. Further, the symbolic representations of fixed point sets for practical systems has been found to be typically very small; the social justification for this observation being that developers rely on simple invariants when designing systems. Thus, the learning based method scales well to real-world systems.

We developed such learning-based algorithms to verify different types of properties for infinite state systems: safety [163, 162], liveness [164], and branching time properties [165]. The ideas have been implemented in a research tool called LeVer [166]. Our experimental analysis on many examples revealed that the approach scales well and outperforms the best known traditional model checker [161]. The results outlined here resulted in the PhD thesis of one student (Abhay Vardhan).

5.2.2 Randomization in verification

Randomization has proved to be an extremely beneficial paradigm in algorithm design and has been extensively used in the last two decades to develop efficient algorithms to solve practical problems in a variety of domains. However, the use of randomization to combat the challenges in algorithmic verification of systems has been largely unexplored.

In this project we developed a statistical approach to verifying probabilistic systems. Such stochastic systems, which explicitly model the probability of random events taking place, define a probability measure on the space of behaviors. Thus, by drawing random samples of executions, the probability space can be estimated, and one can statistically determine the likelihood that a system is correct. The advantage of this approach is that unlike traditional model checking, the algorithm does not need to consider all possible executions of the system. Randomization allows us to ignore executions that may happen, but which are very unlikely to happen. Another advantage is that a formal system model is not needed; one only needs to simulate the system to get sample runs. The disadvantage is that such an algorithm can never guarantee the correctness of a system; there is always a chance that the algorithm got a biased sample and therefore drew an incorrect conclusion. However, the probability of error can always be made as small as one wants by increasing the number of executions sampled. This idea of statistical model checking has been developed to verify safety properties [140] and for liveness properties [141]. This algorithm has also been implemented in the aforementioned tool called Vesta [142]. Our experimental analysis demonstrated the scalability of this approach and the tool was able to analyze systems much faster than traditional model checking algorithms.

Another context in which, we showed the benefits of randomization was in test suite generation [62] for network protocols. The problem, formally called *conformance testing*, involves determining if an unknown implementation is equivalent to a specification, where both are modeled as

finite state Mealy machines by constructing a *test sequence* based on the specification, which is a sequence of inputs that detects all faulty machines. We present a randomized construction of a polynomially long test sequence; no deterministic construction of a short test-suite is known.

In addition, to the broad research themes identified above, as part of this project, we have also developed learning algorithms for stochastic real-time systems [139] and recursive programs [61], constructing test suites for recursive programs [61], verifying network simulation code [145], and verifying stochastic systems in the presence of uncertainties [143].

5.3 Decidability Results on Discrete and Hybrid Systems

As outlined above, the problem of verifying networks of peer-to-peer systems is in general undecidable, because such systems have infinitely many global states. Key features that pose challenges are — multiple concurrently executing agents communicating through messages, requires modelling the unbounded buffer of undelivered messages; agents performing recursive computation, requires modelling the unbounded call stack; agents having discrete and continuous dynamics, require modelling real-valued variables and real-time. We dealt with each of these features individually (and combined in certain ways) to identify key structural properties that can be exploited to yield decidable algorithms.

Systems without continuous dynamics. We first considered non-recursive, distributed, message passing systems. In [157], we observed that an execution consisting of message sends and receives has special algebraic properties that can be exploited to verify such systems against a variety of properties. Next, we studied the impact of recursion by considering models of sequential, recursive software. Such systems have been found to be conveniently modeled by special pushdown models called *visibly pushdown systems*. In [1], we observed that these models can be characterized using special congruence relations on strings. Congruence based characterizations have been known for finite state systems (regular word languages and regular tree languages) for decades. The existence of such a characterization for infinite state systems is surprising. Moreover, this result has important consequences. First, the congruence based characterization can be exploited to minimize system models [61]. Since the time and space requirements of model checking depends on the size of system models, minimization can be used to help scale to large practical systems. Second, algorithms to learn such models can be developed. Finally, by combining developed in these special cases, we presented model checking algorithms for embedded, event-driven, distributed software systems in [138]. These embedded software systems are both concurrent and recursive, but have a key restriction in terms of how recursion and concurrency interact.

STORMED hybrid Systems. Hybrid systems, that have both discrete and continuous dynamics, are notoriously difficult to analyze algorithmically. Systems for which decision algorithms have been developed either have extremely simple continuous dynamics, like system with only clocks variables (timed automata), and with only variables evolving at constant rates (rectangular hybrid automata), or have extremely simple discrete dynamics that require all variables to be reset at each discrete change (o-minimal systems). We observed that if some continuous variable of the systems is guaranteed to evolve monotonically, then this feature can be exploited to yield decision procedures. In a couple of papers [111, 170], we delineated a large class of hybrid systems, termed **STORMED** hybrid systems, having both interesting continuous and discrete dynamics, for which the problem of verifying safety properties can be shown to be decidable.

5.4 Switched Systems

In the paper [82] we studied computational aspects of the problem of stability of switched systems. This paper is concerned with the problem of finding a quadratic common Lyapunov function for

a large family of stable linear systems. We presented gradient iteration algorithms which give deterministic convergence for finite system families and probabilistic convergence for infinite families. Our results and simulations show, in some scenarios, a favorable comparison with more standard techniques based on linear matrix inequalities.

In the paper [173], we presented constructions of a local and global common Lyapunov function for a finite family of pairwise commuting globally asymptotically stable nonlinear systems. The constructions are based on an iterative procedure, which at each step invokes a converse Lyapunov theorem for one of the individual systems. Our results extend a previously available one which relies on exponential stability of the vector fields.

The more recent paper [87] continues to explore the connection between commutation relations and stability of switched systems. We presented a stability criterion for switched nonlinear systems which involves Lie brackets of the individual vector fields but does not require that these vector fields commute. A special case of the main result says that a switched system generated by a pair of globally asymptotically stable nonlinear vector fields whose third-order Lie brackets vanish is globally uniformly asymptotically stable under arbitrary switching. This generalizes a previously known fact for switched linear systems. To prove the result, we considered an optimal control problem which consists in finding the “most unstable” trajectory for an associated control system, and showed that there exists an optimal solution which is bang-bang with a bound on the total number of switches. This property is obtained as a special case of a reachability result by bang-bang controls which is of independent interest. By construction, our criterion also automatically applies to the corresponding relaxed differential inclusion.

We have also studied various types of stochastic stability of switched systems in which the switching is induced by a random process, and for switched systems driven by white noise. Our approach to this problem is inspired by that for the deterministic case: it combines Lyapunov conditions on the individual subsystems with identifying suitable classes of switching signals (which include, but are not limited to, statistically slow-switching processes). Our results on this problem are described in the paper [19] and more publications are forthcoming.

In our recent work [94] we study the problem of establishing stability for hybrid systems through verification of average dwell-time (slow switching) properties. Once one is able to verify that the hybrid system has a sufficiently large average dwell-time, known results can be invoked to prove that it is stable. We introduce a new type of simulation relation for hybrid automata—switching simulation—which allows us to show that the average dwell-time of one automaton is no less than that of another. We show that the question of whether a given hybrid automaton has average dwell-time can be answered by checking a carefully designed invariant or by solving an optimization problem. The invariant-based method is applicable to any hybrid automaton. For suitable classes of automata the invariant in question can be checked automatically. The optimization-based method is applicable to a restricted class of initialized hybrid automata. For this class, a solution of the optimization problem either gives a counterexample execution that violates the average dwell-time property, or it confirms that the automaton indeed satisfies the property. The optimization-based approach is automatic and complements the invariant-based method in the sense that they can be used in combination to find the unknown average dwell-time of a given hybrid automaton.

In the recent paper [171] we study switched systems with external inputs. We prove that a switched nonlinear system has several useful properties of the input-to-state stability (ISS) type under average dwell-time switching signals if each constituent subsystem is ISS. This extends available results for switched linear systems. We apply our result to stabilization of uncertain nonlinear systems via switching supervisory control, and show that the plant states can be kept bounded in the presence of bounded disturbances when the candidate controllers provide ISS properties with respect to the estimation errors.

In another recent paper [172], we address the invertibility problem for switched systems with both inputs and outputs. This is the problem of recovering the switching signal and the input uniquely given an output and an initial state. In the context of hybrid systems, this corresponds

to recovering the discrete state and the input from partial measurements of the continuous state. In solving the invertibility problem, we introduce the concept of singular pairs for **two systems**. We give a necessary and sufficient condition for a switched system to be invertible, which says that the individual subsystems should be invertible and there should be no singular pairs. When the individual subsystems are invertible, we present an algorithm for finding switching signals and inputs that generate a given output in a finite interval when there is a finite number of such switching signals and inputs.

5.5 Markovian Jump Systems: Uniform Performance

In this work we considered Markovian jump linear systems, systems whose parameters jump according to the state transitions of a finite-state Markov chain. These systems model a certain type of hybrid dynamics, and also provide an exact model for situations where feedback loops are subject to random delays. The parameters of these systems are indexed, and the indices are called the system modes. The papers [72, 73] focus on the discrete-time domain, and consider the problems of uniform exponential stability and uniform disturbance attenuation for Markovian jump linear systems. Here uniformity refers to almost sure stability and 2-gain of the sequences of modes, called switching sequences, that are admissible by the automaton which describes the Markov process. We developed semidefinite programming formulations for the solutions to these problems, and their generalizations, without any assumption on the parameters or admissible switching sequences. The work has deep connection to long-standing work on stability of switched systems.

The results show that these formal design problems can be converted to sequences of semidefinite programs, where accuracy is traded off against computational cost. Past work in the literature on these types of problems has yielded results that could not be applied to realistic problems because either they were too computationally demanding or tended to be conservative to the point of providing unacceptable performance.

Related to this work we have also considered linear switched systems, and the three benchmark problems associated with them: stabilization under arbitrary switching sequences, stabilization under a switching path constraint, and construction of stabilizing switching sequences. For discrete-time switched linear systems, control-oriented complete solutions to the first two problems concerning (uniform) stabilization are given in the papers [73, 72] just described; in [74] we solve the third.

We have recently proposed a new output regulation performance criterion [75]. Exact convex conditions for the analysis and synthesis of discrete-time switched linear systems with autonomous switching sequences are obtained, and the formulation is similar to the common receding-horizon control method for standard linear systems. However, in contrast, our technique provides a means of approximating the infinite-horizon LQG performance with guaranteed closed-loop stability. It appears that this work can also be extended to the distributed framework described earlier in this report.

5.6 Certificates for Nonlinear Dynamics

As part of the program we developed powerful new results and theory on providing verification certificates that guarantee certain properties of nonlinear systems. These accomplishments are now described.

5.6.1 Nonlinear system theory

This is the most theoretical component of our work, in which we investigated fundamental structural properties of nonlinear systems with external inputs and/or outputs. These properties are of interest

in their own right, but are also used to support more application-oriented control design and analysis tools, as is clear from the descriptions given below.

At the earlier stages of the project, we were working on understanding the minimum-phase property of nonlinear systems, which is essentially the property that smallness of the output should imply the smallness of the input and the state. Unlike the standard linear notion and its nonlinear analog developed by Isidori and Byrnes in the 1980s, we wanted to formulate a “robust” notion (rather than concentrating on trajectories along which the output is identically zero). Our paper [78] treats such a notion, which we called “output-input stability,” for the general case of multi-input, multi-output nonlinear systems. For systems affine in controls, we derived a necessary and sufficient condition for output-input stability, which relies on a global version of the nonlinear structure algorithm. This condition leads naturally to a globally asymptotically stabilizing state feedback strategy for affine output-input stable systems.

Another fundamental system-theoretic notion that we addressed in our work is that of observability, which is the ability to recover the internal state from output measurements. This property is very well understood for linear systems, but for nonlinear systems this is not the case and there are several possible avenues. In the paper [55] we proposed several definitions of observability for nonlinear systems and explored relationships among them. These observability properties involve the existence of a bound on the norm of the state in terms of the norms of the output and the input on some time interval. A Lyapunov-like sufficient condition for observability was also obtained. As an application, we proved several variants of LaSalle’s stability theorem for switched nonlinear systems. These results were demonstrated to be useful for control design in the presence of switching as well as for developing stability results of Popov type for switched feedback systems.

More recently, we have been looking into disturbance attenuation properties of systems described by nonlinear continuous dynamics and discrete impulses. A desirable response to disturbances was formulated in terms of the input-to-state stability (ISS) property, which was introduced by Sontag in 1989 and has since then become a standard notion in nonlinear system theory. The recent paper [56] introduces appropriate concepts of input-to-state stability (ISS) and integral-ISS for impulsive systems. We provide a set of Lyapunov-based sufficient conditions for establishing these ISS properties. When the continuous dynamics are ISS but the impulses are not, the impulses should not occur too frequently, which can be formalized in terms of an average dwell-time condition. Conversely, when the impulses are ISS but the continuous dynamics are not, there must not be overly long intervals between impulses, which we formalized in terms of a novel reverse average dwell-time condition. We also investigated the cases where both the continuous and discrete dynamics are ISS and when one of these is ISS and the other only marginally stable for the zero input. In the former case we obtained a stronger notion of ISS, for which a necessary and sufficient Lyapunov condition is available. The use of these results was illustrated through examples from a Micro-Electro-Mechanical System (MEMS) oscillator and a problem of remote estimation over a communication network.

5.6.2 Polynomial computation of invariant sets

For nonlinear control systems, one would often like to know the region of attraction of an equilibrium point. Often, this region is difficult to both find and represent computationally. In this program we have developed an approach using polynomials to represent the domain of attraction, and semidefinite programming to perform the computation. The algorithm is iterative, and proceeds by adverting the sublevel set of the polynomial under the inverse flow map.

The paper [174] presents a method for computing the domain of attraction for non-linear dynamical systems. A method is developed where sets are represented as sublevel sets of polynomials. The problem of flowing these sets under the advection map of a dynamical system is converted to a semidefinite program, which is used to compute the coefficients of the polynomials.

The usual mathematical tool used for analysis of the region of attraction is Lyapunov’s method. This gives us a sufficient condition for local stability, although it is often difficult to find a Lyapunov

function that can be used as a certificate for the whole domain of attraction. Several prior approaches have used semidefinite programming to find a quadratic function whose sublevel-set is a good inner approximation to the region of attraction. For system in which the region is complicated, an ellipsoid may not provide a good approximation, and the above methods leave a large unexplored region within the domain of attraction.

With recent developments in algebra and sum-of-squares techniques, it is now possible to solve for a Lyapunov function with a more general polynomial form. Positive definiteness properties are replaced by sum-of-squares constraints which can be efficiently solved using convex optimization. This approach has also allowed finding a Lyapunov function within some specified semi-algebraic region. However, while this provides a method to certify a given inner approximation to the region of attraction, it does not immediately provide a way to find it.

Our research makes use of backward advection of a small initial neighborhood of the equilibrium in order to give an algorithm that in many cases converges to the true domain of attraction. The approach is similar in spirit to the level-set methods that have been used for computation of reachable sets. The key distinction is that most level-set methods represent the function on a mesh; we represent the function as a polynomial. The consequence of this is that the computational requirements may grow more slowly with dimension, if one may fix a-priori the required degrees of the polynomials. By contrast, a mesh-based method has computational costs which grow exponentially with dimension.

Although we do not give the algorithmic details in this report, we show the following numerical example. Consider the Van der Pol oscillator

$$\begin{aligned}\dot{x} &= -y \\ \dot{y} &= x - y(1 - x^2)\end{aligned}$$

The system is locally stable around the origin. We use an initial sublevel set given by the quadratic polynomial $g_0 = 4x^2 + 4y^2 - 1$, which can be verified to be positively invariant, and this is advected with a time step of 0.2. The even-numbered iterates g_0, g_2, g_4, \dots are shown in Figure 9. Some of the iterates are below, normalized to allow integer coefficients.

$$\begin{aligned}p_2 &= -1000 + 2252y^2 - 88y^4 + 11y^6 - 907xy - 56xy^3 - 4xy^5 + 3883x^2 \\ &\quad + 360x^2y^2 - 57x^2y^4 + 660x^3y - x^3y^3 - 417x^4 + 21x^4y^2 + 81x^5y + 260x^6 \\ p_4 &= -1000 + 1614y^2 - 137y^4 + 16y^6 - 1654xy - 170xy^3 + 14xy^5 + 3162x^2 \\ &\quad + 480x^2y^2 - 43x^2y^4 + 94x^3y - 35x^3y^3 + 144x^4 - 2x^4y^2 + 192x^5y + 335x^6 \\ p_{28} &= -10000 + 2510y^2 - 56y^4 + 2y^6 - 4306xy + 42xy^3 + 4xy^5 + 4099x^2 \\ &\quad + 25x^2y^2 + 2x^2y^4 + 1103x^3y - 27x^3y^3 - 687x^4 - x^4y^2 + 2x^5y + 84x^6\end{aligned}$$

It can also be seen that the iterates gradually approach the exact boundary of the domain of attraction. After 30 iterations, the solution covers most of the stable region. After 40 iterations, the stopping criteria allowing an absolute radial change of 0.01 has been met. The final result is shown in Figure 9.

5.6.3 Semialgebraic fundamentals

The work in the previous subsection, and of the earlier subsection on polynomial and semialgebraic games, rely on advances in the understanding of the theory of real polynomials and semialgebraic sets. During the course of this project we have contributed directly to this fundamental theory focusing on geometric aspects [108, 109, 110, 115, 116]. This work brings understanding of when positive semidefinite polynomials can be decomposed as sums-of-squares, and when polynomials are positive on restricted domains. Two conferences ([4, 5]) have been co-organized by one of our PIs on closely related fundamentals.




Figure 9: The Van der Pol oscillator showing the sequence of iterations

6 COMMUNICATIONS FOR NETWORKED CONTROL

A ubiquitous feature of aerial multi-vehicle systems is that they communicate using **wireless links**. The research accomplishments now described are about theory, protocols, and algorithms that are specifically developed to address the special quality-of-service requirements for communication and computing in vehicle — and more generally real-time — systems.

6.1 Control-oriented Information Theory

If all the communication links in a distributed control system are of infinite bandwidth and zero delay then there is no reason to treat the problem as a distributed problem. We can easily construct a new centralized controller with links to all the plants. If the channels are finite bandwidth then thought has to be put into what signals we want to send across them.

These communication links can be noisy, have delays, and drop signals. Furthermore, they may have memory. Thus these communication channels can be considered to be plants themselves. The channel encoders and channel decoders can be considered to be controllers. By viewing channels as plants and encoders and decoders as controllers we are able to unify the different components of the distributed system.

A very important part of this program has been the development of a unified theory of communication and control. In some sense, this requires a unification of the theory of partially-observed stochastic control with information theory. There is a significant difficulty here since the main theorem of Information Theory, the Noisy Channel Coding Theorem, which states that reliable communication in the sense that the probability of error goes to zero, can be achieved if the rate of communication across the channel is less than the capacity and otherwise does not require infinite delay to realize this performance. In problems of control where the system to be controlled is unstable and sensors and actuators are linked to controllers through noisy communication channels, large delays can clearly not be tolerated. In the two part paper [127, 129], a fundamental investigation of these questions has been undertaken. It turns out that Channel Capacity as defined by Shannon does not capture the fundamental limitation of interconnected control and communication systems.

A different notion of reliability—anytime reliability—which captures the fact that the decoder has to act faster than the instability of the system, is needed to provide necessary and sufficient conditions for stabilization to be possible. One of the fundamental conditions of this paper is that the control problem of stabilization is equivalent to a certain communication requirement across the feedback loop.

The problem of stabilization is in some sense a universal problem of control in the same way as digital communication, in the sense of Shannon, is a universal problem. If additional performance requirements are imposed on the control system then control problems will not reduce to communication problems. Energy considerations will come into the picture and the tradeoff between transpiring energy as well as information across the feedback loop needs to be understood. This is a subject of my current research.

In a related paper [153], the analogue of LQG control when the sensor and controller is linked via an Additive White Gaussian Channel is studied. Linear stochastic control systems are examined when there is a communication channel connecting the sensor to the controller. The problem consists of designing the channel encoder and decoder as well as the controller to satisfy some given control objectives. In particular, the role communication has on the classical LQG problem is examined. Conditions under which the classical separation property between estimation and control holds and the certainty equivalent control law is optimal are given. Then the sequential rate distortion framework is presented. The bounds on the achievable performance are presented and the inherent tradeoffs between control and communication costs are shown. In particular, it is shown that optimal quadratic cost decomposes into two terms: a full knowledge cost and a sequential rate distortion cost.

In related work [128], we examine an estimation problem where an unstable source signal is to be estimated by appropriate coding and decoding when the signal is transmitted over a Noisy Channel.

Our understanding of information in systems has been based on the foundation of memoryless processes. Extensions to stable Markov and auto-regressive processes are classical. Berger proved a source coding theorem for the marginally unstable Wiener process, but the infinite-horizon exponentially unstable case had been open since Grays 1970 paper. There were also no theorems showing what is needed to transport such processes across noisy channels.

In this work, we give a fixed rate source coding theorem for the infinite-horizon problem of coding an exponentially unstable Markov process. The encoding naturally results in two distinct bitstreams that have qualitatively different QoS requirements for subsequent transport over a noisy medium. The first stream captures the information that is accumulating within the nonstationary process and requires sufficient anytime reliability on the part of any channel used to transport the process. The second part of the source-code captures the historical information that dissipates within the process and is essentially classical. A converse demonstrating the fundamentally layered nature of such sources is given by means of information-embedding ideas.

There are connections of this work with the new Information and entropy flow picture of Kalman filtering [100] and more generally nonlinear filtering where the filter stores the minimal amount of information necessary to interpret the present and future behavior of the state and dissipates historical information at an optimal rate governed by the Fisher information related to the observations. Contributions also appear in [152, 151].

6.2 Control with Limited Information: Deterministic Systems

We have also been working towards developing a comprehensive theory of nonlinear control with limited information. The type of scenario we have in mind is where the plant and the controller are exchanging information with each other and, due to communication or security constraints, this information is very restricted: coarsely quantized, infrequently updated, delayed, and so on. The main questions then are, how much information is really necessary for control, and what the control law should be (in particular, what robustness properties it should have). Traditional control theory

which assumes perfect and instantaneous signal transmission is inadequate for this task. However, the nonlinear system theory tools that we have been developing can be used to study robustness to errors such as those arising from incomplete information.

In the paper [77], we considered the problem of stabilizing a linear time-invariant system using sampled encoded measurements of its state or output. We derived a relationship between the number of values taken by the encoder and the norm of the transition matrix of the open-loop system over one sampling period, which guarantees that global asymptotic stabilization can be achieved. A coding scheme and a stabilizing control strategy were described explicitly.

In the paper [80], we extended the framework of [77] to nonlinear dynamics. We demonstrated that global asymptotic stabilization is possible if a suitable relationship holds between the number of values taken by the encoder, the sampling period, and a system parameter, provided that a feedback law achieving input-to-state stability (ISS) with respect to measurement errors can be found. The issue of relaxing the latter condition was also studied, and has subsequently led to the work in [56] which we described in Section 5.6.1.

The paper [76] was concerned with global asymptotic stabilization of continuous-time systems subject to dynamic quantization. A hybrid control strategy originating in our earlier work relies on the possibility of making discrete on-line adjustments of quantizer parameters. We explored this method for general nonlinear systems with general types of quantizers affecting the state of the system, the measured output, or the control input. The analysis involves merging tools from Lyapunov stability, hybrid systems, and input-to-state stability.

In [12] state quantization schemes for feedback stabilization of control systems with limited information is investigated, with the focus on designing the least destabilizing quantizer subject to a given information constraint. We explored several ways of measuring the destabilizing effect of a quantizer on the closed-loop system, including (but not limited to) the worst-case quantization error. In each case, we showed how quantizer design can be naturally reduced to a version of the so-called multicenter problem from locational optimization. Algorithms for obtaining solutions to such problems, all in terms of suitable Voronoi quantizers, were discussed. In particular, an iterative solver was developed for a novel weighted multicenter problem which most accurately represents the least destabilizing quantizer design. Simulation studies were also presented.

In the paper [79] we demonstrated that a unified study of quantization and delay effects in nonlinear control systems is possible by merging our quantized feedback control methodology with the small-gain approach to the analysis of functional differential equations with disturbances proposed earlier by Teel. We proved that under the action of a robustly stabilizing feedback controller in the presence of quantization and time delays satisfying suitable conditions, solutions of the closed-loop system starting in a given region remain bounded and eventually enter a smaller region. We presented several versions of this result and showed how it enables global asymptotic stabilization via a dynamic quantization strategy.

In the recent work [81], we consider the problem of achieving input-to-state stability (ISS) with respect to external disturbances for control systems with linear dynamics and quantized state measurements. Quantizers considered in this paper take finitely many values and have an adjustable “zoom” parameter. Building on an approach applied previously to systems with no disturbances in [76], we developed a control methodology that counteracts an unknown disturbance by switching repeatedly between “zooming out” and “zooming in”. Two specific control strategies that yield ISS were presented. The first one is implemented in continuous time and analyzed with the help of a Lyapunov function, similarly to earlier work. The second strategy incorporates time sampling, and its analysis is novel in that it is completely trajectory-based and utilizes a cascade structure of the closed-loop hybrid system. We learned that in the presence of disturbances, time-sampling implementation requires an additional modification which has not been considered in previous work. In [180] input-output stabilization is considered in an ℓ_p context and explicit channel conditions and an algorithm are provided for stabilization.

In [179] decentralized stabilization is considered with finite bandwidth constraints, and sufficient

conditions are provided on the link bandwidths for system stabilization.

6.3 Delay Aware Wireless Networks

Because the cost of even single latencies, if sufficiently large, can be catastrophic in multi-vehicle applications, part of the program has targeted wireless systems that are aware and adapt to transmission delays.

6.3.1 Fundamental delay bounds

In [103] we developed a fundamental tradeoff between network throughput and delay in a mobile wireless network [103]. Using a simple node mobility model and a cell partitioned network structure, we establish that the ratio of delay to throughput must be greater than the number of nodes, N , in the network (i.e., delay/throughput $\geq O(N)$). We also developed algorithms that reduce delays in the network by sending redundant packets along multiple paths. This relatively recent work has already been cited extensively and served as the basis for much research in the field. This work is significant in that it establishes basic performance limits for wireless networks.

6.3.2 Transmission scheduling schemes for time-critical data

In [176] we developed transmission scheduling schemes for meeting deadline constraints over a time varying wireless channel. Such deadline constraints are critical for military command and control communications and are generally difficult to meet in a wireless environment due to channel fluctuations. Our algorithms minimize energy consumption while at the same time meeting the deadline constraints. Using techniques from Dynamic Programming, our algorithms minimize energy consumption for transmitting data with deadline constraints by opportunistically scheduling transmissions at times that the channel is relatively good. Moreover, we developed a novel approach to energy efficient transmission scheduling with general Quality of Service (QoS) requirements [176]. Our approach uses cumulative curves to describe data arrivals, departures, and QoS requirements. Energy efficiency is achieved by spreading the data transmission over time to exploit the convexity of the relationship between power and data rate and further, by opportunistically adapting to the channel variations. We obtain minimum energy transmission policies for meeting a wide range of service requirements, such as delay constraints, buffer limitations, and the transmission of real-time data (e.g., voice or video).

In [177] we formulate the problem of minimizing the energy consumption subject to the deadline constraint as a continuous-time optimal control problem and obtain an analytical solution to the optimal transmission rate. Moreover, using a simple decomposition approach we are able to extend our optimal solution to include the consideration of multiple packet arrivals and variable deadline constraints (i.e., different deadlines for each packet) [178].

6.3.3 Wardrop routing

Routing protocols for multi-hop wireless networks have traditionally used shortest-path routing to obtain paths to destinations, and do not consider traffic load or delay as an explicit factor in the choice of routes. We have formally established that if the number of sources is not too large, then it is possible to construct a perfect flow-avoiding routing, which can boost the throughput provided to each user over that of shortest-path routing by a factor of four when carrier sensing can be disabled, or a factor of 3.2 otherwise [114]. We have also designed a multi-path, load adaptive routing protocol that is generally applicable even when there are more sources. The protocol converges to a Wardrop equilibrium, defined as one where all utilized paths from a source to destination have the same delay, which is less than that over all unutilized paths [112, 113]. We have also addressed the architectural

challenges confronted in the software implementation of a multi-path, delay feedback based, probabilistic routing algorithm. Our routing protocol is (i) completely distributed, (ii) automatically load balances flows, (iii) uses multiple paths whenever beneficial, (iv) guarantees loop-free paths at every time instant even while the algorithm is still converging, and (v) is elegantly implementable in the operating system kernel.

6.4 Synchronization of Clocks in Wireless Systems

In many cooperative missions agents require a common notion of time by which to synchronize events. We have several achievements on creating such a common clock, and an algorithm that to the PIs' knowledge is the best currently available.

Fundamental limitations on clock synchronization in networks We have determined fundamental impossibility results on clock synchronization in wireline or wireless networks, and sharply characterizes what is what is not feasible [45]. Consider a network of n nodes with *affine* clocks, with one node designated as a *reference*. Each other node's clock is described by a *skew* (relative speedup with respect to the reference clock), and an *offset* at time 0 (say) with respect to the reference clock. In order to establish impossibility results, we allow for noiseless communication of messages, that may contain any information that the transmitting node knows about or from current or past packets that it has sent or received. The synchronization problem consists of estimating all the unknown parameters, skews and offsets of all the clocks, as well as the delays of all the communication links. All unknown parameters are assumed to be time-invariant, for sharply delineating impossibility results.

We have proved that the estimation of all unknown parameters is impossible. We show that all nodal skews, as well as all *round-trip delays* between every pair of nodes, can be estimated correctly. However, the vector of unknown link delays and clock offsets can only be determined up to an $(n - 1)$ -dimensional subspace. Each degree of freedom in this subspace corresponds exactly to the offset of one of the $(n - 1)$ clocks with respect to the reference clock. On the positive side, we have shown that every transmitting node can predict precisely the time indicated by the receiver's clock at the instant it receives the packet.

If we further invoke *causality*, that packets cannot be received before they are transmitted, the uncertainty set can be reduced to a polyhedron in R^{n-1} . We have provided necessary and sufficient conditions on the network topology for the polyhedron to be compact and have a non-empty interior.

We have further studied the problem of *receiver-receiver synchronization*, where only receipt times are available, but no time-stamping is done by the sender. We have shown that all nodal skews can still be estimated correctly, but delay differences between neighboring communication links with a common sender can only be characterized up to an affine transformation of the $(n - 1)$ unknown offsets. Moreover, causality does not help; the uncertainty set remains as a translate of R^{n-1} .

We have also investigated structured models for link delays as the sum of a transmitter-dependent delay, a receiver-dependent delay, and a propagation delay, where the latter is known, e.g., via GPS position information. We have identified conditions on the transmission and reception delays which permit a unique solution, and conditions under which the number of the residual degrees of freedom is a constant independent of network size.

Synchronization algorithm We have developed a distributed algorithm to achieve accurate clock synchronization in large multihop wireless networks [146]. The central idea is to exploit the large number of global constraints that have to be satisfied by a common notion of time in a multihop network. If, at a certain time, O_{ij} is the clock offset between two neighboring nodes i and j , then for any loop $i_1, i_2, i_3, \dots, i_n, i_{n+1} = i_1$ in the multihop network, these offsets must satisfy the global constraint $\sum_{k=1}^n O_{i_k i_{k+1}} = 0$. Noisy estimates \hat{O}_{ij} of O_{ij} are usually arrived at by bilateral exchanges

of timestamped messages or local broadcasts. By imposing the large number of global constraints for all the loops in the multihop network, these estimates can be smoothed and made more accurate. We have designed a fully distributed and asynchronous algorithm to exploit all these global constraints. It functions by simple asynchronous broadcasts at each node. Changing the time reference node for synchronization is also easy, consisting simply of one node switching on adaptation, and another switching it off. The algorithm has been implemented on a Berkeley Motes testbed of hundred nodes, and comparative evaluation against a leading algorithm has been performed.

6.5 Performance in Mobile Wireless Systems

Strong research results have been obtained on efficient utilization and deployment of mobile wireless resources.

6.5.1 Optimal flow control scheme for maximizing network throughput and utility

In [104] we developed a novel flow control algorithm for maximizing network utility in heterogeneous networks that include both wireless and wired components. Our algorithm decides when to admit packets into the network based on network layer queue information and does not require knowledge of traffic or channel statistics. We show that when used in conjunction with the routing and power allocation scheme in [105] (also developed under this project) our algorithm maximizes network utility (e.g., throughput); even when the network is overloaded. The above result is significant in that it solves the important problem of optimally controlling a stochastic network when the traffic exceeds the networks capacity. This novel flow control algorithm is very simple to implement in a distributed manner and can be applied to a wide range of commercial and military communication systems such as mobile networks (for command and control) and hybrid networks that include wired and wireless components.

6.5.2 Joint routing and power allocation for wireless networks

In [105] we develop an optimal routing and power allocation strategy for wireless networks with time varying channel conditions and mobile nodes. Our algorithm is optimized across the physical, medium access and network layers. We show that the physical layer power allocation decisions should be made taking network layer queue backlog information. Our routing strategy also makes routing decisions based solely on queue backlog information at the different nodes. In so doing, our optimal routing and power allocation strategy requires minimal signaling overhead (all that needs to be exchanged between nodes and across layer are the queue backlog information). An important feature of our optimal algorithm is that routing and power allocation decisions can be made without knowledge of the traffic statistics or channel statistics; but only based on the queue backlog information.

6.5.3 Randomized algorithms for distributed network control

In [101] we developed a novel framework for distributed scheduling in wireless networks using randomized algorithms. A major challenge in the design of wireless networks is the need for distributed scheduling algorithms that will efficiently share the common channel. Recently, a few distributed scheduling algorithms for networks with different interference constraints have been presented. In networks with primary interference constraints these algorithms guarantee 50% of the maximum possible throughput; and even lower throughput values are achieved under more general interference constraints. In [101] we presented the first distributed scheduling framework that guarantees maximum (100%) throughput. It is based on a combination of a distributed randomized matching algorithm and an algorithm that compares and merges successive matching solutions. We showed that if the matching and compare algorithms satisfy simple conditions related to their performance

and to the inaccuracy of the comparison, the framework attains 100% throughput. We showed that the complexities for achieving 100% throughput are comparable to those of the algorithms that achieve 50% throughput. This work received the Best Paper Award at the ACM SIGMETRICS 2006 conference. In [43] we extended the framework to general interference constraints and in [42] to overall utility maximization (as opposed to throughput).

6.5.4 Maximizing network throughput via partitioning

In [9] we developed a novel partitioning approach for maximizing throughput in a wireless network using distributed scheduling. Our approach is based on a new concept of local pooling whereby networks that satisfy local pooling conditions can achieve 100% throughput using distributed greedy scheduling algorithms. We show that certain classes of graphs, and in particular trees, satisfy this local pooling conditions and are amenable to distributed scheduling. We then partition the network into multiple tree-based subnetworks; where each subnetwork operates on a separate channel. The resulting network, consisting of independent trees can use greedy maximal matching-based scheduling algorithms to achieve high throughput efficiency.

6.5.5 Transmission scheduling for MIMO channels

In [57] we study the problem of efficient scheduling of transmissions to users on a broadcast channel; where the transmitter is equipped with M transmit antennas and each of the receivers is equipped with a single antenna. It is well known that a technique called dirty paper coding, can achieve the capacity of the wireless broadcast channel. However, this requires full knowledge of the channel state information for all users; something that is not practical for most systems of interest. Hence we consider limited feedback schemes whereby we transmit only to a suitably chosen subset of the users. We show that if we only consider a subset of L strong users (i.e., users with high channel gains) and transmit to a subset of M of the L users (where M is the number of antenna elements), then the achieved data rate is asymptotically close to the channel capacity. Hence, we demonstrate that full utilization can be achieved by a simple feedback scheme that does not require knowledge of the channel state of all of the users.

6.6 Increasing Reliability in Cooperative Routing

Cooperative routing takes advantage of the inherent redundancy of wireless networks to increase network reliability [59]. First, the broadcast nature of wireless communications allows nodes to overhear a message that is not intended for them. These nodes can help relay that message to its destination and hence increase network reliability. The receiver, node can combine transmissions from multiple relay nodes to further increase reliability (or alternatively reduce energy consumption). This approach breaks with the traditional layered approach to networking whereby routing is done solely at the network layer. Instead, with this new approach routing is done across both the network and the physical layer. One way to view this approach is as an extension of the advanced physical layer MIMO techniques to the network layer where each node acts as an antenna and several nodes collaborate to provide many of the same benefits including diversity, range extension, and improved link quality to the network layer. In [59] we developed cooperative routing schemes for static wireless networks and demonstrated that approximately 50% energy savings can be achieved via cooperative routing.

We then extend this approach to a wireless fading channel in [60] where we study the problem of communication reliability and diversity in multi-hop wireless networks. To that end, we adopt the outage probability model for a fading channel to develop a probabilistic model for a wireless link. This model establishes a relationship between the link reliability, the distance between communicating nodes and the transmission power. Applying this probabilistic model to a multi-hop network setting, we define and analyze the end-to-end route reliability and develop algorithms for finding the optimal

route between a pair of nodes. The idea of route diversity is introduced as a way to improve the end-to-end route reliability by taking advantage of the wireless broadcast property, the independence of fade state between different pairs of nodes, and the space diversity created by multiple relay nodes along the route. Our results suggest that route diversity can fundamentally change the tradeoff between reliability and power in a multi-hop network.

7 PUBLICATIONS

References

- [1] R. Alur, V. Kumar, P. Madhusudan, and M. Viswanathan. Congruences for Visibly Push-down Languages. In *Proceedings of the International Colloquium on Automata, Languages and Programming*, pages 1102–1114, 2005.
- [2] M. Archer, H. Lim, N. Lynch, S. Mitra, and S. Umeno. Specifying and proving properties of timed I/O automata in the TIOA toolkit. In *Fourth ACM-IEEE International Conference on Formal Methods and Models for Codesign (MEMOCODE’06)*, Napa Valley, California, July 2006.
- [3] A. Arsic and F. Frazzoli. Efficient routing of multiple vehicles with no communications. *Int. J. Robust Nonlinear Control*, 2007. To appear.
- [4] A. Barvinok, M. Beck, C. Haase, B. Reznick, and V. Welker eds. Snowbird ’05 proceedings: Integer points in polyhedra – geometry, number theory, algebra, optimization. *Cont. Math.*, 371, 2005.
- [5] A. Barvinok, M. Beck, C. Haase, B. Reznick, and V. Welker eds. Snowbird ’06 proceedings: Integer points in polyhedra – geometry, number theory, algebra, optimization, statistics. to appear in *Cont. Math.*, 2007.
- [6] D. J. Bertsimas and G. J. van Ryzin. A stochastic and dynamic vehicle routing problem in the Euclidean plane. *Operations Research*, 39:601–615, 1991.
- [7] M. Brown. Air traffic control using virtual stationary automata. Master of Engineering Thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, September 2007.
- [8] M. Brown, S. Gilbert, N. Lynch, C. Newport, T. Nolte, and M. Spindel. The virtual node layer: A programming abstraction for wireless sensor networks. In *Proceedings of the the International Workshop on Sensor Network Architecture (WWSNA)*, April 2007.
- [9] A. Brzezinski, G. Zussman, and E. Modiano. Distributed throughput maximization in wireless mesh networks - a partitioning approach. *Proceedings ACM MOBICOM*, 2006.
- [10] F. Bullo and J. Cortés. Adaptive and distributed coordination algorithms for mobile sensing networks. In V. Kumar, N. E. Leonard, and A. S. Morse, editors, *Cooperative Control. (Proceedings of the 2003 Block Island Workshop on Cooperative Control)*, volume 309 of *Lecture Notes in Control and Information Sciences*, pages 43–62. Springer Verlag, New York, 2005.
- [11] F. Bullo and D. Liberzon. On quantized control and geometric optimization. In *IEEE Conf. on Decision and Control*, pages 2567–2572, Maui, HI, December 2003.
- [12] F. Bullo and D. Liberzon. Quantized control via locational optimization. *IEEE Trans. Automat. Control*, 51:2–13, 2006.

- [13] R. Canetti, L. Cheung, D. Kaynar, M. Liskov, N. Lynch, O. Pereira, and R. Segala. Analyzing security protocols using time-bounded task-PIOAs. *Journal of Discrete Event Dynamic Systems*. To appear.
- [14] R. Canetti, L. Cheung, D. Kaynar, M. Liskov, N. Lynch, O. Pereira, and R. Segala. Using probabilistic I/O automata to analyze an oblivious transfer protocol. ePrint Report 2005/452, Cryptology ePrint Archive, 2005. Also, Technical Report MIT-CSAIL-TR-2007-011, CSAIL, Massachusetts Institute of Technology, Cambridge, MA, February 16, 2007. (Earlier version is MIT-CSAIL-TR-2006-047.
- [15] R. Canetti, L. Cheung, D. Kaynar, M. Liskov, N. Lynch, O. Pereira, and R. Segala. Task-structured probabilistic I/O automata. In *Proceedings the 8th International Workshop on Discrete Event Systems (WODES 2006)*, Ann Arbor, Michigan, July 2006.
- [16] R. Canetti, L. Cheung, D. Kaynar, N. Lynch, and O. Pereira. Modeling bounded computation in long-lived systems. In progress.
- [17] R. Canetti, L. Cheung, D. Kaynar, N. Lynch, and O. Pereira. Compositional security for task-pioas. In *Proceedings of the 20th IEEE Computer Security Foundations Symposium*, Venice, Italy, July 2007.
- [18] R. Chadha and M. Viswanathan. Decidability results for well-structured transition systems with auxiliary storage. In *Proceedings of the Conference on Concurrency Theory*, pages 136–150, 2007.
- [19] D. Chatterjee and D. Liberzon. Stability analysis of deterministic and stochastic switched systems via a comparison principle and multiple lyapunov functions. *SIAM J. Control Optim.*, 15:174–206, 2006.
- [20] L. Cheung, N. Lynch, R. Segala, and F. Vaandrager. Switched probabilistic I/O automata. In Z. Liu and K. Araki, editors, *Proceedings of the First International Colloquium on Theoretical Aspects of Computing (ICTAC2004)*, Guiyang, China, September 2004, volume 3407 of *Lecture Notes in Computer Science*, pages 494–510. Springer, 2006. Long version in Nijmegen Institute for Computing and Information Sciences (NIII) Technical Report NIII-R0437, Catholic University of Nijmegen, Nijmegen, The Netherlands, September 2004.
- [21] L. Cheung, N. Lynch, R. Segala, and F. Vaandrager. Switched probabilistic PIOA: Parallel composition via distributed scheduling. *Theoretical Computer Science*, 365(1-2):83–108, November 2006.
- [22] E. M. Clarke, O. Grumberg, and D. A. Peled. *Model Checking*. MIT Press, 2000.
- [23] R. Cogill and S. Lall. Control design for topology-independent stability of interconnected systems. In *Proceedings of the American Control Conference*, pages 3717–3722, June 2004.
- [24] R. Cogill and S. Lall. Decentralized stochastic decision problems and polynomial optimization. In *Proceedings of Allerton Conference on Communication, Control, and Computing*, pages 1477–1485, 2001.
- [25] R. Cogill and S. Lall. Topology independent controller design for networked systems. In *Proceedings of the IEEE Conference on Decision and Control*, pages 1788–1793, 2004.
- [26] R. Cogill and S. Lall. Decentralized stochastic decision problems and polynomial optimization. In *Proceedings of the American Control Conference*, pages 2709–2714, 2005.

- [27] R. Cogill and S. Lall. An approximation algorithm for the discrete team decision problem. *SIAM Journal on Control and Optimization*, 45(4):1359–1368, 2006.
- [28] R. Cogill and S. Lall. Backlog analysis of maximal matching switching with speedup. In *Proceedings of the IEEE Conference on Decision and Control*, pages 1876–1881, 2006.
- [29] R. Cogill and S. Lall. Suboptimality bounds in stochastic control: A queueing example. In *Proceedings of the American Control Conference*, pages 1642–1647, 2006.
- [30] R. Cogill, S. Lall, and J. Hespanha. A constant factor approximation algorithm for event-based sampling. In *Proceedings of the American Control Conference*, pages 305–311, 2007.
- [31] R. Cogill, S. Lall, and P. A. Parrilo. On structured semidefinite programs for the control of symmetric systems. In *Proceedings of the Allerton Conference on Communication, Control, and Computing*, pages 1536–1545, 2003.
- [32] R. Cogill, M. Rotkowitz, B. Van Roy, and S. Lall. An approximate dynamic programming approach to decentralized control of stochastic systems. In *Proceedings of Allerton Conference on Communication, Control, and Computing*, pages 1040–1049, 2004.
- [33] J. Cortés and F. Bullo. From geometric optimization and nonsmooth analysis to distributed coordination algorithms. In *IEEE Conf. on Decision and Control*, pages 3274–3280, Maui, HI, December 2003.
- [34] J. Cortés and F. Bullo. Coordination and geometric optimization via distributed dynamical systems. *SIAM Journal on Control and Optimization*, 44(5):1543–1574, 2005.
- [35] J. Cortés, S. Martínez, T. Karatas, and F. Bullo. Coverage control for mobile sensing networks. *IEEE Transactions on Robotics and Automation*, 20(2):243–255, 2004.
- [36] M. Demirbas, A. Arora, T. Nolte, and N. Lynch. Brief announcement: Stalk— a self-stabilizing hierarchical tracking service for sensor networks. In *Proceedings of the 23rd Annual ACM Symposium on Principles of Distributed Computing (PODC 2004)*, page 378, St. John’s, Newfoundland, Canada, July 2004.
- [37] Z. Di, M. Farhood, and G.E. Dullerud. Control of distributed systems over graphs. In *Proceedings of IEEE Conference on Decision and Control*, 2006.
- [38] G. E. Dullerud and R. D’Andrea. Distributed control of heterogeneous systems. *IEEE Transactions on Automatic Control*, 49:2113–2128, 2004.
- [39] J. J. Enright and E. Frazzoli. UAV routing in a stochastic time-varying environment. In *IFAC World Congress*, Prague, Czech Republic, July 2005. Electronic Proceedings.
- [40] J. J. Enright and E. Frazzoli. The stochastic traveling salesman problem for the Reeds-Shepp car and the differential drive robot. In *IEEE Conf. on Decision and Control*. 2006.
- [41] J. J. Enright, E. Frazzoli, K. Savla, and F. Bullo. On multiple UAV routing with stochastic targets: performance bounds and algorithms. In *AIAA Conf. on Guidance, Navigation and Control*, August 2005. Electronic Proceedings.
- [42] A. Eryilmaz, A. Ozdaglar, and E. Modiano. Randomized algorithms for throughput-optimality and fairness in wireless networks. *Proceedings IEEE Conference on Decision and Control (CDC)*, 2006.
- [43] A. Eryilmaz, A. Ozdaglar, and E. Modiano. Polynomial complexity algorithms for full utilization of multi-hop wireless networks. *Proceedings IEEE Infocom*, 2007.

- [44] E. Frazzoli and F. Bullo. Decentralized algorithms for vehicle routing in a stochastic time-varying environment. In *IEEE Conf. on Decision and Control*, pages 3357–3363, Paradise Island, Bahamas, December 2004.
- [45] Nikolaos M. Freris, Scott R. Graham, and P. R. Kumar. Fundamental limits on synchronizing clocks over networks. *submitted to IEEE Transactions on Automatic Control*, 2007.
- [46] A. Ganguli, J. Cortés, and F. Bullo. On rendezvous for visually-guided agents in a nonconvex polygon. In *IEEE Conf. on Decision and Control and European Control Conference*, pages 5686–5691, Seville, Spain, December 2005.
- [47] A. Ganguli, J. Cortés, and F. Bullo. Distributed deployment of asynchronous guards in art galleries. In *American Control Conference*, pages 1416–1421, Minneapolis, MN, June 2006.
- [48] A. Ganguli, J. Cortés, and F. Bullo. Maximizing visibility in nonconvex polygons: Nonsmooth analysis and gradient algorithm design. *SIAM Journal on Control and Optimization*, 45(5):1657–1679, 2006.
- [49] A. Ganguli, J. Cortés, and F. Bullo. Distributed coverage of nonconvex environments. In V. Saligrama, editor, *Networked Sensing Information and Control (Proceedings of the NSF Workshop on Future Directions in Systems Research for Networked Sensing, May 2006, Boston, MA)*, Lecture Notes in Control and Information Sciences, pages 289–305. Springer Verlag, 2007.
- [50] A. Ganguli, J. Cortés, and F. Bullo. Multirobot rendezvous with visibility sensors in nonconvex environments. *IEEE Transactions on Robotics*, August 2007. (Submitted Nov 2006) Conditionally Accepted.
- [51] A. Ganguli, J. Cortés, and F. Bullo. Visibility-based multi-agent deployment in orthogonal environments. In *American Control Conference*, pages 3426–3431, New York, July 2007.
- [52] A. Ganguli, S. Susca, S. Martínez, F. Bullo, and J. Cortés. On collective motion in sensor networks: Sample problems and distributed algorithms. In *IEEE Conf. on Decision and Control and European Control Conference*, pages 4239–4244, Seville, Spain, December 2005.
- [53] S. Garland, N. Lynch, J. Tauber, and M. Vaziri. IOA user guide and reference manual. Technical Report MIT-LCS-TR-961, MIT Computer Science and Artificial Intelligence Laboratory, Cambridge, MA 02139, July 2004.
- [54] C. Georgiou, P. Mavrommatis, and J. A. Tauber. Implementing asynchronous distributed systems using the ioa toolkit. Technical Report MIT-LCS-TR-966, MIT CSAIL, Cambridge, MA, September 2004.
- [55] J. P. Hespanha, D. Liberzon, D. Angeli, and E. D. Sontag. Nonlinear observability notions and stability of switched systems. *IEEE Trans. Automat. Control*, 50:154–168, 2005.
- [56] J. P. Hespanha, D. Liberzon, and A. R. Teel. On input-to-state stability of impulsive systems. *Automatica*. To appear.
- [57] K. Jagannathan, S. Borst, P. Whiting, and E. Modiano. Scheduling of multi-antenna broadcast systems with heterogeneous users. *IEEE Journal of Selected Areas in Communications*, 2007.
- [58] D. K. Kaynar, N. Lynch, R. Segala, and F. Vaandrager. *Theory of Timed I/O Automata*. Synthesis Lectures on Computer Science. Morgan-Claypool Publishers, May 2006. Also, revised and shortened version of Technical Report MIT-LCS-TR-917a (from 2004), MIT Laboratory for Computer Science, Cambridge, MA.

- [59] A. Khandani, J. Aboundi, E. Modiano, and L. Zhang. Cooperative routing in wireless networks. *IEEE Transactions on Communications*, 2007.
- [60] A. Khandani, E. Modiano, J. Abounadi, and L. Zheng. Reliability and route diversity in wireless networks. *Proceedings of Conference on Information Science and System*, 2005.
- [61] V. Kumar, P. Madhusudan, and M. Viswanathan. Minimization, Learning and Conformance Testing of Boolean Programs. In *Proceedings of the Conference on Concurrency Theory*, pages 203–217, 2006.
- [62] V. Kumar and M. Viswanathan. Conformance Testing in the Presence of Multiple Faults. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, pages 1136–1145, 2005.
- [63] J. Le Ny, M. Dahleh, and E. Feron. Dynamic programming and linear programming relaxations for bandit problems with switching costs. submitted.
- [64] J. Le Ny, M. Dahleh, and E. Feron. Multi-agent task assignment in the bandit framework. In *Proceedings of the 45th IEEE Conference on Decision and Control*, San Diego, CA, December 2006.
- [65] J. Le Ny, M. Dahleh, and E. Feron. Multi-uav dynamic routing with partial observations using restless bandits allocation indices. Technical Report 2776, Massachusetts Institute of Technology, 2007. Submitted to the American Control Conference 2008.
- [66] J. Le Ny and E. Feron. An approximation algorithm for the curvature-constrained traveling salesman problem. In *Proceedings of the 43rd Annual Allerton Conference on Communications, Control and Computing*, September 2005.
- [67] J. Le Ny and E. Feron. Restless bandits with switching costs: Linear programming relaxations, performance bounds and limited lookahead policies. In *Proceedings of the American Control Conference*, Minneapolis, MN, June 2006.
- [68] J. Le Ny and E. Feron. Performance evaluation of a multi-agent risk-sensitive tracking system. In *Proceedings of the 46th IEEE Conference on Decision and Control*, 2007.
- [69] J. Le Ny, E. Frazzoli, and E. Feron. Complexity and approximation algorithms for the Dubins traveling salesman problem. in preparation.
- [70] J. Le Ny, E. Frazzoli, and E. Feron. The curvature-constrained traveling salesman problem for high point densities. In *Proceedings of the 46th IEEE Conference on Decision and Control*, 2007.
- [71] J. Lee and G.E. Dullerud. A dynamic decentralized sequential multi-hypothesis testing problem under uniformly distributed nonstationary observations. In *Proceedings of the IEEE Conference on Decision and Control*, 2004. journal version in review.
- [72] J.-W. Lee and G. E. Dullerud. Optimal disturbance attenuation for discrete-time switched and markovian jump linear systems. *SIAM Journal of Control and Optimization*, 45:1329–1358, 2006.
- [73] J.-W. Lee and G. E. Dullerud. A stability and contractiveness analysis of discrete-time markovian jump linear systems. *Automatica*, 42:205–218, 2006.
- [74] J.-W. Lee and G. E. Dullerud. Uniformly stabilizing sets of switching sequences for switched linear systems. *IEEE Transactions on Automatic Control*, 52:868–874, 2007.

- [75] J.-W. Lee, G. E. Dullerud, and P.P. Khargonekar. An output regulation problem for switched linear systems in discrete time. In *Proceedings of IEEE Conference on Decision and Control*, 2007.
- [76] D. Liberzon. Hybrid feedback stabilization of systems with quantized signals. *Automatica*, 39:1513–1554, 2003.
- [77] D. Liberzon. On stabilization of linear systems with limited information. *IEEE Trans. Automat. Control*, 48:304–307, 2003.
- [78] D. Liberzon. Output-input stability implies feedback stabilization. *Systems Control Lett.*, 53:237–248, 2004.
- [79] D. Liberzon. Quantization, time delays, and nonlinear stabilization. *IEEE Trans. Automat. Control*, 51:1190–1195, 2006.
- [80] D. Liberzon and J. P. Hespanha. Stabilization of nonlinear systems with limited information feedback. *IEEE Trans. Automat. Control*, 50:910–915, 2005.
- [81] D. Liberzon and D. Nešić. Input-to-state stabilization of linear systems with quantized state measurements. *IEEE Trans. Automat. Control*, 52:767–781, 2007.
- [82] D. Liberzon and R. Tempo. Switched systems, common Lyapunov functions, and gradient algorithms. *IEEE Trans. Automat. Control*, 49:990–994, 2004.
- [83] N. Lynch, S. Mitra, and T. Nolte. Motion coordination using virtual nodes. In *Forty-Fourth IEEE Conference on Decision and Control and European Control Conference (CDC-ECC 2005)*, Seville, Spain, December 2005. Invited paper. Also, Technical Report MIT-LCS-TR-986, MIT CSAIL, Cambridge, MA, April 2005.
- [84] N. Lynch, R. Segala, and F. Vaandraager. Hybrid I/O automata. *Information and Computation*, 185(1):105–157, August 2003. Also, Technical Report MIT-LCS-TR-827d, MIT Laboratory for Computer Science, Cambridge, MA 02139, January 13, 2003.
- [85] N. Lynch, R. Segala, and F. Vaandrager. Compositionality for probabilistic automata. In Roberto Amadio and Denis Lugiez, editors, *Concurrency Theory: The 14th International Conference on Concurrency Theory (CONCUR 2003)*, Marseille, France, September 2003, volume 2761 of *Lecture Notes in Computer Science*, pages 208–221. Springer, 2003.
- [86] N. Lynch, R. Segala, and F. Vaandrager. Observing branching structure through probabilistic contexts. *SIAM Journal on Computing*, 37(4):977–1013, 2006.
- [87] M. Margaliot and D. Liberzon. Lie-algebraic stability conditions for nonlinear switched systems and differential inclusions. *Systems Control Lett.*, 55:8–16, 2006.
- [88] S. Martínez, F. Bullo, J. Cortés, and E. Frazzoli. On synchronous robotic networks – Part I: Models, tasks and complexity. *IEEE Trans. Automat. Control*, 52(12):2199–2213, 2007.
- [89] S. Martínez, F. Bullo, J. Cortés, and E. Frazzoli. On synchronous robotic networks – Part II: Time complexity of rendezvous and deployment algorithms. *IEEE Trans. Automat. Control*, 52(12):2214–2226, 2007.
- [90] J.C. Mitchell, A. Ramanathan, A. Scedrov, and V. Teague. A probabilistic polynomial-time process calculus for the analysis of cryptographic protocols. *Theoretical Computer Science*, 353(1):118–164, 2006.

- [91] S. Mitra and M. Archer. Reusable pvs proof strategies for proving abstraction properties of i/o automata. In *STRATEGIES 2004: 5th International Workshop on Strategies in Automated Deduction*, Cork, Ireland, July 2004.
- [92] S. Mitra and M. Archer. PVS strategies for proving abstraction properties automata. *Electronic Notes in Theoretical Computer Science*, 125(2):45–65, 2005.
- [93] S. Mitra and D. Liberzon. Stability of hybrid automata with average dwell time: an invariant approach. In *Proceedings of the 43rd Conference on Decision and Control (CDC 2004)*, Paradise Island, Bahamas, December 2004.
- [94] S. Mitra, D. Liberzon, and N. Lynch. Verifying average dwell time of hybrid systems. *ACM Transactions in Embedded Computing Systems*. To appear.
- [95] S. Mitra and N. Lynch. Approximate simulations for task-structured probabilistic I/O automata. In *LICS workshop on Probabilistic Automata and Logics (PAul06)*, Seattle, WA, August 2006.
- [96] S. Mitra and N. Lynch. Proving approximate implementation relations for probabilistic I/O automata. *ENTCS*, 174(8):71–93, 2007.
- [97] S. Mitra and N. Lynch. Trace-based semantics of probabilistic timed I/O automata. In Alberto Bemporad, Antonio Bicchi, and Giorgio C. Buttazzo, editors, *Hybrid Systems: Computation and Control (HSCC 2007, Pisa, Italy, April, 2007)*, volume 4416 of *Lecture Notes in Computer Science*, pages 718–722. Springer, 2007.
- [98] S. Mitra, Y. Wang, N. Lynch, and E. Feron. Application of Hybrid I/O automata in safety verification of pitch controller for model helicopter system. Technical Report MIT-LCS-TR-880. Laboratory for Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139, January 2003.
- [99] S. Mitra, Y. Wang, N. Lynch, and E. Feron. Safety verification of model helicopter controller using Hybrid Input/Output automata. In O. Maler and A. Pnueli, editors, *Hybrid Systems: Computation and Control. 6th International Workshop (HSCC 2003), Prague, the Czech Republic April 3-5, 2003*, volume 2623 of *Lecture Notes in Computer Science*, pages 343–358. Springer-Verlag, 2003. Also, long version in Technical Report MIT-LCS-TR-880, MIT Laboratory for Computer Science, Cambridge, MA, January 2003.
- [100] S.K. Mitter and N.J. Newton. Information and entropy flow in the kalman-bucy filter. *J. of Stat. Phys.*, 118:145–176.
- [101] E. Modiano, D. Shah, and G. Zussman. Maximizing throughput in wireless networks via gossiping. *Proceedings ACM SIGMETRICS / IFIP Performance*, 2006. winner of best paper award).
- [102] P. A. Parrilo N. Stein and A. Ozdaglar. Characterization and computation of correlated equilibria in infinite games. In *Proceedings IEEE Conference on Decision and Control*, 2007.
- [103] M. Neely and E. Modiano. Capacity and delay tradeoffs for ad hoc mobile networks. *IEEE Transactions on Information Theory*, 2005.
- [104] M. Neely and E. Modiano. Fairness and optimal stochastic control for heterogeneous networks. *IEEE Infocom*, 2005. Extended version to appear in *IEEE/ACM Transactions on Networking* in 2008.

- [105] M. Neely, E. Modiano, and C. Rohrs. Dynamic power allocation and routing for time-varying wireless networks. *IEEE Journal of Selected Areas in Communications*, 2005.
- [106] K. J. Obermeyer, A. Ganguli, and F. Bullo. Asynchronous distributed searchlight scheduling. In *IEEE Conf. on Decision and Control*, pages 4863–4868, New Orleans, LA, December 2007.
- [107] P.A. Parrilo. Polynomial games and sum of squares optimization. In *Proceedings IEEE Conference on Decision and Control*, 2006.
- [108] V. Powers and B. Reznick. Polynomials positive on unbounded rectangles. *Positive polynomials in Control* (A. Garulli, D. Henrion eds., Springer), *Lecture Notes in Control and Inform. Sci.*, 312:151–163, 2005.
- [109] V. Powers and B. Reznick. A quantitative polya’s theorem with corner zeros. *Proceedings of the 2006 International Symposium of Symbolic and Algebraic Computation*, J.-G. Dumas editor, New York, ACM Press, pages 285–290, 2006.
- [110] V. Powers, B. Reznick, C. Scheiderer, and Sottile F. A new approach to hilbert’s theorem on ternary quartics. *C. R. Acad. Sci. Paris*, 339:617–620, 2004.
- [111] P. Prabhakar, V. Vladimerou, M. Viswanathan, and G. Dullerud. A decidable class of planar linear hybrid systems. In *Proceedings of the Conference on Hybrid Systems Control and Computation*, 2008, to appear.
- [112] Vivek Raghunathan and P. R. Kumar. A wardrop routing protocol for ad hoc wireless networks. In *Proceedings of the IEEE Conference on Decision and Control*, pages 4661–4666, 2004.
- [113] Vivek Raghunathan and P. R. Kumar. Issues in wardrop routing in wireless networks. In *Proceedings of WICON’05, First International Wireless Internet Conference*, pages 34–41, 2005.
- [114] Vivek Raghunathan and P. R. Kumar. Wardrop routing in wireless networks: from theory to implementation. *submitted to IEEE Transactions on Mobile Computing*, 2007.
- [115] B. Reznick. Clean lattice tetrahedra. submitted for publication.
- [116] B. Reznick. On hilbert’s construction of positive polynomials. submitted for publication.
- [117] C.M. Robson. TIOA and UPPAAL. Master of Engineering Thesis, MIT Department of Electrical Engineering and Computer Science, Cambridge, MA, May 2004.
- [118] M. Rotkowitz. *Tractable Problems in Optimal Decentralized Control*. PhD thesis, Stanford University, June 2005.
- [119] M. Rotkowitz, R. Cogill, and S. Lall. A simple condition for the convexity of optimal control over networks with delays. In *Proceedings of the IEEE Conference on Decision and Control*, pages 6686–6691, 2005.
- [120] M. Rotkowitz and S. Lall. Decentralized control information structures preserved under feedback. In *Proceedings of the IEEE Conference on Decision and Control*, pages 569–575, December 2002.
- [121] M. Rotkowitz and S. Lall. Decentralized control of unstable systems and quadratically invariant information constraints. In *Proceedings of the IEEE Conference on Decision and Control*, pages 2865–2871, December 2003.

- [122] M. Rotkowitz and S. Lall. Decentralized control subject to communication and propagation delays. In *Proceedings of the IEEE Conference on Decision and Control*, pages 4990–4997, 2001.
- [123] M. Rotkowitz and S. Lall. On computation of optimal controllers subject to quadratically invariant sparsity constraints. In *Proceedings of the American Control Conference*, pages 5659–5664, June 2004.
- [124] M. Rotkowitz and S. Lall. Affine controller parameterization for decentralized control over banach spaces. *IEEE Transactions on Automatic Control*, 51(9):1497–1500, September 2006.
- [125] M. Rotkowitz and S. Lall. A characterization of convex problems in decentralized control. *IEEE Transactions on Automatic Control*, 51(2):274–286, 2006.
- [126] M. Rotkowitz and S. Lall. Convexification of optimal decentralized control without a stabilizing controller. In *Proceedings of the International Symposium on Mathematical Theory of Networks and Systems (MTNS)*, pages 1496–1499, 2006.
- [127] A. Sahai and S.K. Mitter. The necessity and sufficiency of anytime capacity for stabilization of a linear system over a noisy communication link, part i: Scalar systems. *IEEE Trans. on Information Theory*, 52(8):3369–3395, 2006.
- [128] A. Sahai and S.K. Mitter. Source coding and channel requirements for unstable processes. *submitted to IEEE Trans. on Information Theory*, 2007.
- [129] A. Sahai and S.K. Mitter. The necessity and sufficiency of anytime capacity for stabilization of a linear system over a noisy communication link, part ii: Vector systems. *to appear in IEEE Trans. on Information Theory*, 2008.
- [130] K. Savla. *Multi UAV Systems with Motion and Communication Constraints*. PhD thesis, Electrical and Computer Engineering Department, University of California at Santa Barbara, Santa Barbara, August 2007. Available electronically at <http://ccdc.mee.ucsb.edu>.
- [131] K. Savla, F. Bullo, and E. Frazzoli. On traveling salesperson problems for Dubins’ vehicle: stochastic and dynamic environments. In *IEEE Conf. on Decision and Control and European Control Conference*, pages 4530–4535, Seville, Spain, December 2005.
- [132] K. Savla, F. Bullo, and E. Frazzoli. On Traveling Salesperson Problems for a double integrator. In *IEEE Conf. on Decision and Control*, pages 5305–5310, San Diego, CA, December 2006.
- [133] K. Savla, F. Bullo, and E. Frazzoli. The coverage problem for loitering Dubins vehicles. In *IEEE Conf. on Decision and Control*, pages 1398–1403, New Orleans, LA, December 2007.
- [134] K. Savla, F. Bullo, and E. Frazzoli. Traveling Salesperson Problems for a double integrator. *IEEE Trans. Automat. Control*, 2007. (Submitted Nov 2006) To appear.
- [135] K. Savla, E. Frazzoli, and F. Bullo. On the point-to-point and traveling salesperson problems for Dubins’ vehicle. In *American Control Conference*, pages 786–791, Portland, OR, June 2005.
- [136] K. Savla, E. Frazzoli, and F. Bullo. On the Dubins Traveling Salesperson Problems: Novel approximation algorithms. In G. S. Sukhatme, S. Schaal, W. Burgard, and D. Fox, editors, *Robotics: Science and Systems II (Proceedings of the Second RSS Conference, August 2006, Philadelphia PA)*. MIT Press, Cambridge, MA, 2007. Available electronically at <http://arxiv.org/abs/cs.R0/0603010>.

- [137] K. Savla, E. Frazzoli, and F. Bullo. Traveling Salesperson Problems for the Dubins vehicle. *IEEE Trans. Automat. Control*, 53(9), 2008. (Submitted Jun 2006) To appear.
- [138] K. Sen and M. Viswanathan. Model checking multi-threaded programs with asynchronous atomic methods. In *Proceedings of the Conference on Computer-Aided Verification*, pages 300–314, 2006.
- [139] K. Sen, M. Viswanathan, and G. Agha. Learning Continuous Time Markov Chains. In *Proceedings of the International Conference on Quantitative Evaluation of Systems*, pages 116–155, 2004.
- [140] K. Sen, M. Viswanathan, and G. Agha. Statistical Model Checking of Black-Box Probabilistic Systems. In *Proceedings of the Conference on Computer-Aided Verification*, 2004.
- [141] K. Sen, M. Viswanathan, and G. Agha. Stochastic Model Checking of Probabilistic Systems. In *Proceedings of the Conference on Computer-Aided Verification*, 2005.
- [142] K. Sen, M. Viswanathan, and G. Agha. VESTA: A Statistical Model-checker and Analyzer for Probabilistic Systems. In *Proceedings of the International Conference on Quantitative Evaluation of Systems*, pages 251–252, 2005.
- [143] K. Sen, M. Viswanathan, and G. Agha. Model checking markov chains in the presence of uncertainties. In *Proceedings of the International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 394–410, 2006.
- [144] P. Shah and P. A. Parrilo. Polynomial stochastic games via sum of squares optimization. In *Proceedings IEEE Conference on Decision and Control*, 2007.
- [145] A. Sobeih, M. Viswanathan, and J. Hou. Check and Simulate: A Case for Incorporating Model Checking in Network Simulation. In *Proceedings of ACM-IEEE International Conference on Formal Methods and Models for Codesign*, pages 27–36, 2004.
- [146] R. Solis, V.S. Borkar, and P. R. Kumar. A new distributed time synchronization protocol for multihop wireless networks. In *Proceedings of the IEEE Conference on Decision and Control*, pages 2734–2739, 2006.
- [147] E. Solovey. Simulation of composite I/O automata. Master of Engineering Thesis, MIT Department of Electrical Engineering and Computer Science, Cambridge, MA, August 2003.
- [148] J. M. Steele. Probabilistic and worst case analyses of classical problems of combinatorial optimization in Euclidean space. *Mathematics of Operations Research*, 15(4):749, 1990.
- [149] N. Stein, A. Ozdaglar, and P. A. Parrilo. Separable and low-rank continuous games. In *Proceedings IEEE Conference on Decision and Control*, 2006.
- [150] A. Stubbs, V. Vladimerou, A. Fulford, J. Strick, and G. E. Dullerud. A hovercraft testbed for networked and decentralized control. *IEEE Control Systems Magazine*, 26:56–69, 2006.
- [151] S. Tatikonda and S.K. Mitter. Control over noisy channels. *IEEE Trans. on Automatic Control*, 49(7):1196–1201, 2004.
- [152] S. Tatikonda and S.K. Mitter. Control under communication constraints. *IEEE Trans. on Automatic Control*, 49(7):1056–1068, 2004.
- [153] S. Tatikonda, A. Sahai, and S.K. Mitter. Stochastic linear control over a communication channel. *IEEE Trans. on Automatic Control (Special Issue on Networked Control Systems)*, 49(9):1549–1561, 2004.

- [154] J. A. Tauber. *Verifiable Compilation of I/O Automata without Global Synchronization*. PhD thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139, February 2005.
- [155] J.A. Tauber and S.J. Garland. Definition and expansion of composite automata in IOA. Technical Report MIT/LCS/TR-959, Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA, July 2004.
- [156] J.A. Tauber, N.A. Lynch, and M.J. Tsai. Compiling ioa without global synchronization. In *Proceedings of the 3rd IEEE International Symposium on Network Computing and Applications (IEEE NCA04)*, Cambridge, MA, August 2004.
- [157] P. Thati and M. Viswanathan. Verification of Asynchronous Systems with Unbounded and Unordered Message Buffers. In *Proceedings of the Workshop on Automated Verification of Infinite-State Systems*, 2004.
- [158] S. Umeno. Proving safety properties of an aircraft landing protocol using timed and untimed I/O automata: a case study. Master's thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, October 2006.
- [159] S. Umeno and N. Lynch. Safety verification of an aircraft landing protocol: A refinement approach. In *Hybrid Systems: Computation and Control (HSCC 2007)*, Pisa, Italy, April 3-5, 2007, volume 4416 of *Lecture Notes in Computer Science*, pages 557–572.
- [160] S. Umeno and N. Lynch. Proving safety properties of an aircraft landing protocol using I/O automata and the PVS theorem prover: a case study. In *Formal Methods, International Symposium of Formal Methods Europe (FM 2006)*, Lecture Notes in Computer Science, Hamilton, Ontario Canada, August 2006. Springer.
- [161] A. Vardhan. *Learning to Verify Systems*. PhD thesis, University of Illinois, Urbana-Champaign, 2005.
- [162] A. Vardhan, K. Sen, M. Viswanathan, and G. Agha. Actively Learning to Verify Safety for FIFO Automata. In *Proceedings of the International Conference on Foundations of Software Technology and Theoretical Computer Science*, pages 494–505, 2004.
- [163] A. Vardhan, K. Sen, M. Viswanathan, and G. Agha. Learning to Verify Safety Properties. In *Proceedings of the International Conference on Formal Engineering Methods*, pages 274–289, 2004.
- [164] A. Vardhan, K. Sen, M. Viswanathan, and G. Agha. Using Language Inference to Verify ω -regular Properties. In *Proceedings of the International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 45–60, 2005.
- [165] A. Vardhan and M. Viswanathan. Learning to Verify Branching-time Properties. In *Proceedings of IEEE-ACM International Conference on Automated Software Engineering*, pages 325–328, 2005.
- [166] A. Vardhan and M. Viswanathan. Lever: A tool for learning based verification. In *Proceedings of the Conference on Computer-Aided Verification*, pages 471–474, 2006.
- [167] M. Vaziri, J.A. Tauber, M.J. Tsai, and N. Lynch. Systematic removal of nondeterminism for code generation in I/O automata. Technical Report MIT/LCS/TR-960, Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA, July 2004.

- [168] M. Viswanathan and R. Viswanathan. Foundations for Circular Compositional Reasoning. In *Proceedings of the International Colloquium on Automata, Languages and Programming*, pages 835–847, 2001.
- [169] M. Viswanathan and R. Viswanathan. A Higher Order Modal Fixed Point Logic. In *Proceedings of the Conference on Concurrency Theory*, 2004.
- [170] V. Vladimirov, P. Prabhakar, M. Viswanathan, and G. Dullerud. STORMED Hybrid Systems. Unpublished manuscript, 2007.
- [171] L. Vu, D. Chatterjee, and D. Liberzon. Input-to-state stability of switched systems and switching adaptive control. *Automatica*, 43:639–646, 2007.
- [172] L. Vu and D. Liberzon. Invertibility of switched linear systems. *Automatica*. To appear.
- [173] L. Vu and D. Liberzon. Common Lyapunov functions for families of commuting nonlinear systems. *Systems Control Lett.*, 54:405–416, 2005.
- [174] T. Wang, S. Lall, and M. West. Polynomial level-set methods for nonlinear dynamical systems analysis. In *Proceedings of the Allerton Conference on Communication, Control, and Computing*, pages 640–649, 2005.
- [175] P. Whittle. Restless bandits: activity allocation in a changing world. *Journal of Applied Probability*, 25A:287–298, 1988.
- [176] M. Zafer and E. Modiano. A calculus approach to minimum energy transmission policies with quality of service guarantees. *IEEE Infocon*, 2005.
- [177] M. Zafer and E. Modiano. Delay constrained energy efficient data transmission over a wireless fading channel. *Workshop on Information Theory and Application, University of California, San Diego, February*, 2007.
- [178] M. Zafer and E. Modiano. Minimum energy transmission over a wireless fading channel with packet deadlines. *Proceedings of IEEE Conference on Decision and Control (CDC)*, 2007.
- [179] C. Zhang and G.E. Dullerud. Decentralized control with bandwidth constraints. In *Proceedings of the American Control Conference*, 2007.
- [180] C. Zhang and G.E. Dullerud. Finite gain stabilization with logarithmic quantization. In *Proceedings of the IEEE Conference on Decision and Control*, 2007.

ACKNOWLEDGEMENT/DISCLAIMER

This work was sponsored (in part) by the Air Force Office of Scientific Research, USAF, under grant/contract number F49620-02-1-0325. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Office of Scientific Research or the U.S. Government.

8 APPENDICES

A Awards and Honors

Fraucesco Bullo

- *Best Student Paper Prize* (Ganguli, Cortes, Bullo), American Control Conference, 2006.
- *Plenary speaker*, Hybrid Systems: Computation and Control Conference, 2006.
- *ONR Young Investigator Prize* 2003.

Geir E. Dullerud, *Fellow*, Institute of Electrical and Electronics Engineers (IEEE), 2008.

Emilio Frazzoli, *NSF CAREER Award*, 2002.

P.R. Kumar

- *Fellow*, National Academy of Engineering, 2007.
- *Plenary speaker*, Hybrid Systems: Computation and Control Conference, 2006.
- *IEEE Control Systems Field Award*, 2005.

Sanjay Lall

- *George Axelby Best Paper Award*, IEEE Transactions on Automatic Control, 2007.
- *Best Student Paper Prize* (Rotkowitz, Cogil, Lall), IEEE CDC, 2005.

Daniel Liberzon

- Donald P. Eckman Award, American Automatic Control Council, 2007.
- *Young Author Prize*, IFAC World Congress, Barcelona, Spain, 2002.

Nancy A. Lynch

- Van Wijngaarden Prize, 2006.
- Fellow, *National Academy of Engineering* (elected prior to MURI).

Sanjoy Mitter

- Control Heritage Award, *IEEE Control Systems Society*, 2007.
- *Penner Lecture*, University of California, San Diego, 2006.
- Member of *Venice Institute of Arts, Humanities and Science*, 2003.
- Fellow, *National Academy of Engineering* (elected prior to MURI).

Eytan Modiano, *Best Paper Award*, ACM SIGMETRICS, 2006.

Pablo Parrilo

- SIAM Control and Systems Theory Prize, 2005
- Donald P. Eckman Award, American Automatic Control Council, 2005.

B Transitions

Northrup-Grumman. In this application, real-time path planning algorithms based on mixed-integer programming have been adopted by Northrup-Grumman corporation to perform the guidance of small-unmanned rotorcraft UAVs; Firescout. POC: Dr. Robert Miller, robert.h.miller@ngc.com.

Navy Research Laboratories. Developed capability improvements to a front-end, simulator, and translator for PVS theorem-prover at NRL. POC: Myla Archer, tame@itd.nrl.navy.mil.

Honeywell. Development of sum-of-squares based (sos) methods for the safety verification of an Advanced Life Support System (HSCC 2005). Dr. Sonja Glavaski, Sonja.Glavaski@honeywell.com.

VeroModo Inc. Company has been founded to create commercial software, based on Input-Output Automata framework supported by this project, that the military, industry, and academia can use for modelling, analyzing and simulating timing-based systems.
POC: Nancy Lynch, lynch@theory.csail.mit.edu.

Wright-Patterson AFB, AFRL. Solving the Dubins traveling sales person problem with camera pointing constraints and wind disturbances. Systematic prosecution of targets. POC: Dr. Corrie Schumacher, Corey.Schumacher@wpafb.af.mil.

Nascent Technologies/Aurora Flight Sciences. Path planning for single and multiple unmanned vehicles. Vehicle trajectory planning in constrained environments. POC: Dr. James Paduano, paduano@nascent.com.

C Personnel Supported (list cumulative)

Jinane Abounadi, Research Scientist, Electrical Engineering and Computer Science, MIT.

Erin Aylward, SM student, Electrical Engineering and Computer Science, MIT.

Francesco Bullo, Associate Professor, Mechanical and Environmental Engineering, University of California, Santa Barbara.

Debasish Chatterjee, PhD student, Electrical and Computer Engineering, University of Illinois.

Been-Der Chen, PhD student, Aeronautics and Astronautics, Stanford University.

Randy Cogill, PhD student, Electrical Engineering, Stanford University.

Jorge Cortes, Postdoctoral Fellow, Coordinated Science Laboratory, University of Illinois.

Zhe Di, PhD student, Mechanical Engineering, University of Illinois.

Hau Duang, PhD student, Mathematics, University of Illinois.

Geir Dullerud, Professor, Mechanical Engineering, University of Illinois.

John Enright, PhD student, Aerospace Engineering, University of California, Los Angeles.

Eric Feron, Professor, Aerospace Engineering, Georgia Institute of Technology.

Emilio Frazzoli, Associate Professor, Aeronautics and Astronautics, MIT.

Anurag Ganguli, PhD student, Electrical and Computer Engineering, University of Illinois.
 Arvind Giridhar, PhD student, Electrical and Computer Engineering, University of Illinois.
 Lt. Col Scott Graham, PhD student, Electrical and Computer Engineering, University of Illinois.
 Chung H. Hsieh, PhD student, Aerospace Engineering, University of California, Los Angeles.
 Amir Khandani, PhD student, Electrical Engineering and Computer Science, MIT.
 Colleen Kilker, PhD student, Mathematics, University of Illinois.
 Dilsun Kirli, Postdoctoral Fellow, Electrical Engineering and Computer Science, MIT.
 P.R. Kumar, Professor, Electrical and Computer Engineering, University of Illinois.
 Viraj Kumar, PhD student, Computer Science, University of Illinois.
 Sanjay Lall, Associate Professor, Aeronautics and Astronautics, Stanford University.
 Ji-Woong Lee, Postdoctoral Fellow, Coordinated Science Laboratory, University of Illinois.
 Daniel Liberzon, Assistant Professor, Electrical and Computer Engineering, University of Illinois.
 Nancy A. Lynch, Professor, Electrical Engineering and Computer Science, MIT.
 John C. Mitchell, Professor, Computer Science, Stanford University.
 Sayan Mitra, PhD student, Electrical Engineering and Computer Science, MIT.
 Sanjoy Mitter, Professor, Electrical Engineering and Computer Science, MIT.
 Eytan Modiano, Associate Professor, Aeronautics and Astronautics, MIT.
 Jerome Le Nuy, PhD student, Aeronautics and Astronautics, MIT.
 Pablo Parrilo, Associate Professor, Electrical Engineering and Computer Science, MIT.
 Ajith Ramanathan, PhD student, Computer Science, Stanford University.
 Bruce Reznick, Professor, Mathematics, University of Illinois.
 Melissa Simmons, PhD student, Mathematics, University of Illinois.
 Karin Sigurd, PhD student, Electrical Engineering and Computer Science, MIT.
 Noah Stein, PhD student, Electrical Engineering and Computer Science, MIT.
 Andrew Stubbs, PhD student, Mechanical Engineering, University of Illinois.
 Abhay Vardhan, PhD student, Computer Science, University of Illinois.
 Mahesh Viswanathan, Associate Professor, Computer Science, University of Illinois.
 Vladimeros Vladimerou, Electrical and Computer Engineering, University of Illinois.
 Travis Wendt, Graduate student, Aerospace Engineering, University of Illinois.
 Murtaza Zafer, PhD student, Electrical Engineering and Computer Science, MIT.

INSTRUCTIONS FOR COMPLETING SF 298

1. REPORT DATE. Full publication date, including day, month, if available. Must cite at least the year and be Year 2000 compliant, e.g. 30-06-1998; xx-06-1998; xx-xx-1998.

2. REPORT TYPE. State the type of report, such as final, technical, interim, memorandum, master's thesis, progress, quarterly, research, special, group study, etc.

3. DATES COVERED. Indicate the time during which the work was performed and the report was written, e.g., Jun 1997 - Jun 1998; 1-10 Jun 1996; May - Nov 1998; Nov 1998.

4. TITLE. Enter title and subtitle with volume number and part number, if applicable. On classified documents, enter the title classification in parentheses.

5a. CONTRACT NUMBER. Enter all contract numbers as they appear in the report, e.g. F33615-86-C-5169.

5b. GRANT NUMBER. Enter all grant numbers as they appear in the report, e.g. AFOSR-82-1234.

5c. PROGRAM ELEMENT NUMBER. Enter all program element numbers as they appear in the report, e.g. 61101A.

5d. PROJECT NUMBER. Enter all project numbers as they appear in the report, e.g. 1F665702D1257; ILIR.

5e. TASK NUMBER. Enter all task numbers as they appear in the report, e.g. 05; RF0330201; T4112.

5f. WORK UNIT NUMBER. Enter all work unit numbers as they appear in the report, e.g. 001; AFAPL30480105.

6. AUTHOR(S). Enter name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. The form of entry is the last name, first name, middle initial, and additional qualifiers separated by commas, e.g. Smith, Richard, J, Jr.

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES). Self-explanatory.

8. PERFORMING ORGANIZATION REPORT NUMBER. Enter all unique alphanumeric report numbers assigned by the performing organization, e.g. BRL-1234; AFWL-TR-85-4017-Vol-21-PT-2.

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES). Enter the name and address of the organization(s) financially responsible for and monitoring the work.

10. SPONSOR/MONITOR'S ACRONYM(S). Enter, if available, e.g. BRL, ARDEC, NADC.

11. SPONSOR/MONITOR'S REPORT NUMBER(S). Enter report number as assigned by the sponsoring/monitoring agency, if available, e.g. BRL-TR-829; -215.

12. DISTRIBUTION/AVAILABILITY STATEMENT. Use agency-mandated availability statements to indicate the public availability or distribution limitations of the report. If additional limitations/ restrictions or special markings are indicated, follow agency authorization procedures, e.g. RD/FRD, PROPIN, ITAR, etc. Include copyright information.

13. SUPPLEMENTARY NOTES. Enter information not included elsewhere such as: prepared in cooperation with; translation of; report supersedes; old edition number, etc.

14. ABSTRACT. A brief (approximately 200 words) factual summary of the most significant information.

15. SUBJECT TERMS. Key words or phrases identifying major concepts in the report.

16. SECURITY CLASSIFICATION. Enter security classification in accordance with security classification regulations, e.g. U, C, S, etc. If this form contains classified information, stamp classification level on the top and bottom of this page.

17. LIMITATION OF ABSTRACT. This block must be completed to assign a distribution limitation to the abstract. Enter UU (Unclassified Unlimited) or SAR (Same as Report). An entry in this block is necessary if the abstract is to be limited.